



Grid refinement in ICON v2.6.4

Günther Zängl, Daniel Reinert, and Florian Prill

Deutscher Wetterdienst, Offenbach am Main, Germany

Correspondence: Günther Zängl (guenther.zaengl@dwd.de)

Received: 25 April 2022 – Discussion started: 23 May 2022

Revised: 1 September 2022 – Accepted: 5 September 2022 – Published: 23 September 2022

Abstract. This article describes the implementation of grid refinement in the ICOSahedral Nonhydrostatic (ICON) modeling system. It basically follows the classical two-way nesting approach known from widely used mesoscale models like MM5 or WRF, but it differs in the way feedback from fine grids to coarser grids is applied. Moreover, the ICON implementation supports vertical nesting in the sense that the upper boundary of a nested domain may be lower than that of its parent domain. Compared to the well-established implementations on quadrilateral grids, new methods had to be developed for interpolating the lateral boundary conditions from the parent domain to the child domain(s). These are based on radial basis functions (RBFs) and partly apply direct reconstruction of the prognostic variables at the required grid points, whereas gradient-based extrapolation from parent to child grid points is used in other cases. The runtime flow control is written such that limited-area domains can be processed identically to nested domains except for the lateral boundary data supply. To demonstrate the functionality and quality of the grid nesting in ICON, idealized tests based on the Jablonowski–Williamson test case (Jablonowski and Williamson, 2006) and the Schär mountain wave test case (Schär et al., 2002) are presented. The results show that the numerical disturbances induced at the nest boundaries are small enough to be negligible for real applications. This is confirmed by experiments closely following the configuration used for operational numerical weather prediction at DWD, which demonstrate that a regional refinement over Europe has a significant positive impact on the forecast quality in the Northern Hemisphere.

1 Introduction

The ICON (ICOSahedral Nonhydrostatic) modeling framework is jointly developed by the German Weather Service (DWD), the Max Planck Institute for Meteorology (MPI-M), the German Climate Computing Center (DKRZ) and the Karlsruhe Institute for Technology (KIT), targeting a unified global numerical weather prediction (NWP) and climate modeling system (GCM). The development work started in 2004 with basic research on the model grid and the numerical formulation of the dynamical core in a highly idealized shallow-water framework (Bonaventura and Ringler, 2005; Rípodas et al., 2009). The next step was the implementation of a hydrostatic dynamical core using the same model equations, time integration scheme and vertical discretization as the spectral transform ECHAM model (Wan et al., 2013). This work focused on investigating the effects of the different horizontal grid formulation (icosahedral vs. spherical harmonics) and provided the basis for a first version with full physics coupling. In parallel, nonhydrostatic dynamical cores were developed on hexagonal grids (Gassmann and Herzog, 2008; Gassmann, 2013) and triangular grids (Zängl et al., 2015). As discussed in Gassmann (2011) and Danilov (2012), C-grid-type triangular discretizations suffer from a spurious computational mode, which manifests itself in a rapidly oscillating checkerboard pattern in the horizontal divergence field, giving rise to numerical stability problems. To date, this problem is lacking a rigorous mathematical solution. However, Zängl et al. (2015) showed that the problem could largely be mitigated by a specific averaging of the velocity components entering the divergence operator. This pragmatic solution allowed retaining the triangular C-grid discretization. Moreover, a triangular grid proved to be more suitable for implementing a two-way grid nesting capability than a hexagonal one. Triangular cells can be recursively

partitioned into successively smaller triangles, which leads to a unique relationship between parent and child cells. For hexagons, however, this is not the case, as the majority of child cells are shared between two adjacent parent cells. To our knowledge, this is one aspect that makes two-way grid refinement approaches developed for hexagonal grids significantly more complex (e.g., Dubos and Kevlahan, 2013).

Despite impressive advances in computational power over the last decades, the application of global models with uniform, convection-permitting resolution on weather or even climate timescales is still too costly to be performed on a regular basis. To date, high-resolution limited-area models (LAMs) serve as a cost-effective alternative for exploring the mesoscale and microscale, and they will continue to serve as a working horse for both the numerical weather prediction (NWP) and climate community. Limited-area models have proven successful, but they are known to have conceptual deficiencies, such as the potential ill-posedness of lateral boundary conditions (Davies, 2014), possible inconsistencies with the driving model in terms of the governing equations, numerical formulations or physical parameterizations, and the lack of regional- to global-scale interactions (Warner et al., 1997). In order to mitigate these deficiencies, a number of methods have been considered to achieve locally enhanced resolution in global models.

The oldest approach dating back more than 40 years is the usage of stretched grids (Schmidt, 1977; Staniforth and Mitchell, 1978). The so-called Schmidt transformation allows for enhanced resolution in a particular region of interest by redistributing the grid points of an initially uniform model grid. This approach has proven successful in various NWP and climate applications. A review of GCM applications is provided, e.g., by Fox-Rabinovitz et al. (2008), and Goto et al. (2015) discuss a more recent NWP application using NICAM, which is based on a modified Schmidt transformation on an icosahedral grid (Tomita, 2008). Grid stretching obviates the need for lateral boundary conditions, and it allows for an immediate interaction between global and regional scales. Probably the largest drawback of this method is the fact that grid points can only be redistributed rather than created. Enhancing the resolution in one region of the globe implies a coarsening in another region. Besides the inevitable coarsening itself, the insufficient resolution of disturbances passing through the coarsened region may also negatively affect the simulation in the refined region.

More recently, global models using locally refined unstructured meshes have been developed. Unlike the grid stretching approach, they allow new grid cells to be added in regions of interest (static h-refinement). This approach is pursued, for example, by the Model for Prediction Across Scales (MPAS) (Skamarock et al., 2012) and the spectral element dynamical core of the Community Atmosphere Model (CAM) (Zarzycki et al., 2014). Albeit being much more flexible, this approach faces similar challenges as the grid stretching approach mentioned previously. In both approaches, the

time step is restricted by the smallest cell in the domain unless specific measures like sub-stepping are taken for individual cells or regions or unless horizontally implicit numerical methods are chosen. Moreover, care must be taken that the parameterizations are applicable on a wide range of scales and turn off gradually when the respective processes become resolved on the model grid (such as the convection parameterization in the gray zone). The development of scale-aware parameterizations poses major challenges and is an area of active research (Gross et al., 2018). However, notable progress in terms of scale awareness has been reported, e.g., for the CAM5 parameterization suite when compared to CAM4 (Gettelman et al., 2018).

The approach we are pursuing closely resembles traditional two-way nesting, as known from many regional mesoscale models such as MM5 (Grell et al., 1994) or WRF (Skamarock et al., 2019). Two-way nesting differs from the previous single-grid approaches by the fact that multiple grids of different resolution are overlaid onto each other such that individual points on the globe are covered by more than one prognostic grid cell. Scale interaction can be ensured by feeding the nested-grid solution back to the underlying parent grid at regular time intervals. Similar to LAMs, this method requires lateral boundary conditions to be specified for every nested domain, and it may suffer from spurious wave reflections at nest boundaries due to the abrupt jump in resolution. On the other hand, it allows many model settings to be chosen individually for each domain, which improves numerical efficiency and reduces the requirements concerning the parameterization suite. Distinct time steps can be chosen for each domain, allowing, e.g., a proportional reduction in refined domains in order to meet stability constraints. The parameterizations can be tuned individually for each domain and resolution or even switched off, relaxing the requirement of scale awareness to some degree. While two-way nesting has been successfully applied in limited-area modeling for decades, it is still much less common in global modeling. Focusing on recent global models which are being used in research or operational forecasting, the authors are only aware of the two-way nesting approach implemented in the Geophysical Fluid Dynamics Laboratory (GFDL) Finite-Volume Cubed-Sphere Dynamical Core (FV3) (Harris and Lin, 2013, 2014; Mouallem et al., 2022).

The purpose of this article is to describe the implementation of grid nesting in ICON. Compared to previous two-way nesting approaches, it is the first implementation on triangular grids and differs in the way feedback from fine grids to coarser grids is realized. In addition, the ICON implementation allows for vertical nesting in the sense that the upper boundary of a nested domain may be lower than that of the parent domain, which is not supported by most previous nesting implementations in other models. Without specific discussion, we note that a limited-area mode is available in ICON as a by-product of the grid nesting implementation, differing from nesting only in the way the lateral boundary

conditions are provided. A detailed description the grid nesting implementation will be provided in Sect. 2, followed by application examples in Sect. 3. A brief summary will be given in Sect. 4.

2 Domain nesting in ICON

2.1 Description of the basic design

The understanding of ICON's nesting implementation requires some basic knowledge of ICON's mathematical–physical design. We therefore start by summarizing key elements of the dynamical core, time stepping and physics–dynamics coupling scheme. ICON's dynamical core solves the fully compressible, nonhydrostatic Euler equations on the sphere using either the shallow or deep atmosphere formulation (Borchert et al., 2019). The spatial discretization is performed on an unstructured icosahedral–triangular Arakawa–C grid in the horizontal and a terrain-following height-based SLEVE coordinate (Leuenberger et al., 2010) with Lorenz-type staggering in the vertical. As described in Zängl et al. (2015), the prognostic variables encompass the edge-normal horizontal wind speed v_n , vertical wind speed w , total air density ρ , virtual potential temperature θ_v and mass fractions q_k of various moisture quantities. See Fig. 1 in Wan et al. (2013) for the variable placement on the triangular grid. The two-time-level predictor–corrector time integration scheme is fully explicit except for the terms describing the vertical propagation of sound waves, which are treated implicitly (Zängl et al., 2015). Hence, the permitted integration time step is comparatively small and constrained by the ratio of the speed of sound to the horizontal mesh size.

To optimize computational efficiency, different integration time steps are used for the dynamical core on the one hand and additional sub-grid physical processes (and tracer transport) on the other hand. The time steps will be denoted by $\Delta\tau$ for the dynamical core and Δt for physics. The dynamical core is sub-stepped with respect to physical processes, with a single physics time step consisting of five dynamics sub-steps by default ($n_{\text{subs}} = \Delta t / \Delta\tau = 5$).

The physics–dynamics coupling scheme in ICON further distinguishes between *fast* physical processes, having a timescale shorter than or comparable to the time step Δt , and *slow* processes, which are considered to have a large timescale compared to Δt . Processes that fall into the category of fast are saturation adjustment, grid-scale microphysics or turbulent diffusion, while examples of slow processes are convection and radiation. Fast processes are integrated with the previously defined time step $\Delta\tau$, whereas slow processes may be integrated with process-specific larger time steps that are an integer multiple of Δt . We therefore call the time step Δt the fast physics time step. As the tracer transport is performed at Δt as well, we apply the coupling of nested domains at Δt for both dynamics and tracer variables

because this simplifies maintaining consistency with continuity (Gross et al., 2002). Essentially it is required to provide time-averaged mass fluxes (averaged over the sub-steps) for tracer transport at nest lateral boundaries (see Sect. 2.2.1) and to ensure that the child-to-parent feedback increments for partial densities ρq_i sum up to the feedback increments for total density ρ (see Sect. 2.2.2).

The static mesh refinement in ICON is accomplished using multiple individual grids, whereby one or more higher-resolution (child) domains are overlaid on a coarser base (parent) domain. The base domain can be a regional or global domain. Model integration on the child domain is performed in addition to that on the underlying part of the parent domain; i.e., there is no “hole” in the parent domain where the child domain is located.

Each child domain has a defined parent domain providing lateral boundary conditions, but a parent domain can have several child domains. The child domains can be located in different geographical regions and can also serve as parent domains for further subdomains, but domains having the same parent are not allowed to share the same parent grid cells because this would lead to ambiguities in combination with two-way nesting. Nested domains may also be switched on or off during runtime. Conceptually, the number of nested domains is arbitrary and controlled by the grid files provided as input, but of course not all choices would make sense from a physical point of view. Each domain can be regarded as separate instances of the same model that are coupled to each other using the same numerical operators and filters, time integration scheme, and physics–dynamics coupling. If desired, however, different physical settings can be chosen individually for each domain. For example, radiation can be called more frequently on subdomains, or the convection scheme can be tuned differently or even switched off completely.

The refinement ratio between the parent domain and a child domain is fixed to a value of 2, as each parent triangle is split into four child triangles (see the right part of Fig. 1). While higher refinement ratios would technically be possible, we decided for this restriction because it reduces the risk of numerical artifacts (e.g., by partial wave reflection) along nest boundaries. Consistent with the refinement ratio of 2, the model integration time step Δt is multiplied by a factor of 0.5 for each additional nesting level. The coupling time step between successive nesting levels is the fast physics time step Δt .

The parent–child coupling can be either one-way or two-way. A mixture of one-way and two-way coupled domains is also possible. Two-way versus one-way coupling means that the prognostic variables on the child domain are transferred back to the coarser parent domain at regular time intervals using a dedicated feedback mechanism. As a result, the solution on the parent domain benefits from the higher resolution of the child domain. In the case of one-way nesting, the feedback is switched off.

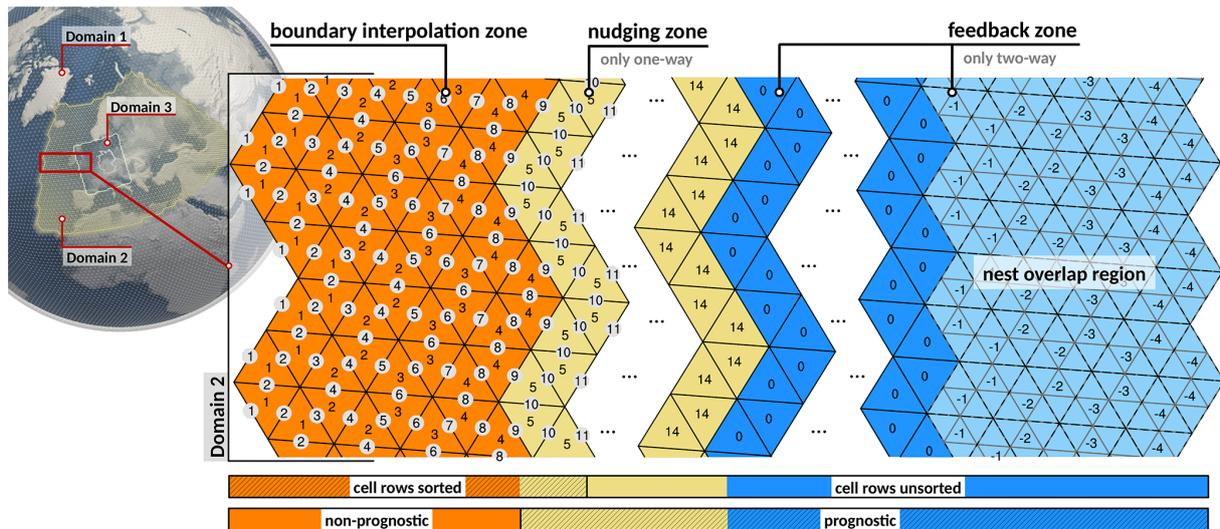


Figure 1. Basic structure of a nested (or limited-area) domain, exemplified by a section of domain 2 shown in the upper left. Orange: boundary interpolation zone, having a fixed width of four cell rows. Ocher: nudging zone with adjustable width that is only active for one-way nesting and in limited-area mode. Blue and light blue: child-to-parent feedback zone. Light blue: nest overlap region, i.e., a region for which a higher-resolution child domain exists (see domain 3 in the schematic in the upper left). Prognostic computations are restricted to the feedback and nudging zone. Integers indicate the internal indexing of domain 2, which is used to assign cells and edges to individual zones. More details on the indexing and the indicated sorting of cell rows are given in Appendices A1 and A2.

To perform the coupling, we conceptually split any nested domain into three zones, which we call the *boundary interpolation zone*, the *nudging zone* and the *feedback zone*. In order to identify grid points belonging to these zones, cells, edges and vertices are indexed according to their distance from the boundary (see Appendix A1 for details). These zones along with the grid point indexing are depicted in Fig. 1, which shows part of the boundary region and inner region of a nested domain. Limited-area domains are treated technically like one-way nested domains except for the fact that the boundary interpolation zone is filled with external input data rather than data interpolated from the parent domain.

The boundary interpolation zone is non-prognostic and is meant to store the boundary conditions that are necessary to solve the governing equations in the child domain. Boundary conditions are needed for the prognostic variables v_n , w , ρ , θ_v and q_k . By a dedicated boundary update mechanism (see Sect. 2.2.1 for details), both the prognostic variables and their time tendencies are interpolated from the parent to the child domain, and the boundary conditions are updated at every child time step. The boundary interpolation zone has a fixed width of four cell rows. This is motivated by the technical constraint that the boundary zone needs to match parent cell rows (i.e., an odd number of cell rows is not allowed), combined with the fact that two cell rows would not be sufficient to cover all stencil operations performed in the dynamical core. For example, the ∇^4 -diffusion operator (Zängl et al., 2015) requires information from three adjacent cell rows. We note that halving the size of the boundary interpolation zone would be possible by modifying some stencil operations near

the lateral boundary, such as replacing ∇^4 with ∇^2 . However, this would have no impact on computational efficiency in practical applications with MPI domain decomposition (see Appendix A2).

The nudging zone, which is active in the case of one-way nesting only, serves to damp differences between the driving solution in the adjacent boundary interpolation zone and the prognostic solution in the child domain. Essentially, the prognostic model state of the child domain is relaxed (nudged) to the parent state following the traditional Newtonian relaxation approach described by Davies (1976). Details of the implementation are provided in Sect. 2.2.3. For two-way nesting, no nudging is applied, and the boundary interpolation zone borders the feedback zone.

In the feedback zone, the model state on the parent domain is relaxed towards the updated model state on the child domain at every fast physics time step Δt_p of the parent domain. We refer to this as *relaxation-type* feedback. As a result, the parent and child domains remain tightly coupled, and the solution on the parent domain benefits from the enhanced resolution of the child domain. Feedback is applied to the prognostic variables v_n , w , ρ and θ_v as well as to the mass fractions of water vapor q_v , cloud water q_c and cloud ice q_i . Child-to-parent feedback has already been successfully applied in mesoscale models like MM5 (Grell et al., 1994) or WRF (Skamarock et al., 2019), as well as in global simulations based on a cubed-sphere grid (Harris and Lin, 2013). However, our approach differs in the sense that the parent state is relaxed towards the child state with an adjustable timescale, rather than being overwritten by the child

state. Compared to the conventional direct feedback, which is available as an option in ICON as well, the relaxation feedback has the advantage of generating fewer numerical disturbances near vertical nest interfaces, leading to slightly better forecast quality in NWP applications (without vertical nesting, the quality difference is small). In addition, the relaxation feedback requires no adjustment of the model orography at the parent grid level, which is more convenient if nested domains are turned on or off during runtime. See Sect. 2.2.2 for further details.

2.2 Parent–child coupling

2.2.1 Lateral boundary update: parent → child

The boundary update mechanism provides the child domain with up-to-date lateral boundary conditions for the prognostic variables v_n , w , ρ , θ_v and q_k . In order to prevent parent-to-child interpolated values of ρ from entering the solution of the mass continuity equation, the above set of variables is extended by the horizontal mass flux ρv_n . This will allow for parent–child mass flux consistency, as described below. For the subsequent description of the algorithm, let the model state on the parent and child domain be denoted by \mathcal{M}_p^n and \mathcal{M}_c^n , respectively, where n specifies the time step index.

In general, the boundary update works as follows. Let ψ_p^n and ψ_p^{n+1} represent a prognostic variable on the parent domain at time steps n and $n + 1$, respectively. Once the variables on the parent domain \mathcal{M}_p have been updated from n to $n + 1$, the time tendency

$$\frac{\partial \psi_p}{\partial t} = \frac{\psi_p^{n+1} - \psi_p^n}{\Delta t_p}$$

is diagnosed. Both the field ψ_p^n at time level n and the tendency $\frac{\partial \psi_p}{\partial t}$ are then interpolated (downscaled) from the parent grid cells or edges to the corresponding cells or edges of the child domain’s boundary zone (orange cells in Fig. 1). With $\mathcal{I}_{p \rightarrow c}$ denoting the interpolation operator, we get

$$\begin{aligned} \psi_c^n &= \mathcal{I}_{p \rightarrow c}(\psi_p^n) \\ \frac{\partial \psi_c}{\partial t} &= \mathcal{I}_{p \rightarrow c}\left(\frac{\partial \psi_p}{\partial t}\right). \end{aligned}$$

The interpolated tendencies are needed in order to provide the lateral boundary conditions at the right time levels, since two integration steps are necessary on the child domain in order to reach the model state \mathcal{M}_c^{n+1} , with each step consisting of several (typically five) dynamics sub-steps. The temporal update is performed at each dynamics sub-step $\Delta \tau_c$ for the prognostic variables of the dynamical core and at each large step Δt_c for the tracer variables. As an example, the boundary conditions at the first dynamics sub-step of the first and second (fast physics) integration step on the child domain read ψ_c^n and $\psi_c^n + 0.5 \Delta \tau_p \partial \psi_c / \partial t$, respectively.

Concerning the interpolation operator $\mathcal{I}_{p \rightarrow c}$, we distinguish between cell-based variables (i.e., scalars) and edge-based variables (v_n and ρv_n). For cell-based variables, a 2D horizontal gradient is reconstructed at the parent cell circumcenter by first computing edge-normal gradients at edge midpoints, followed by a nine-point reconstruction of the 2D gradient at the cell center based on radial basis functions (Narcowich and Ward, 1994). The interpolated value at the j th child cell center is then calculated as

$$\psi_{c_j} = \psi_p + \nabla \psi_p \cdot \mathbf{d}(p, c_j), \quad \forall j \in \{1..4\}, \quad (1)$$

with $\nabla \psi_p$ denoting the horizontal gradient at the parent cell center and $\mathbf{d}(p, c_j)$ the distance vector between the parent and j th child cell center. The same operator is applied to cell-based tendencies.

To prevent excessive overshoots and undershoots of ψ_{c_j} in the vicinity of strong gradients, a limiter for $\nabla \psi_p$ is implemented. It ensures that

$$\frac{1}{\beta} \psi_{p,\min} < \psi_{c_j} < \beta \psi_{p,\max} \quad \forall j \in \{1..4\}$$

on all four child points, where $\psi_{p,\min}$ and $\psi_{p,\max}$ denote the minimum and maximum of ψ_p , respectively, on the above-mentioned reconstruction stencil including the local cell center, and $\beta = 1.05$ is a tuning parameter. To minimize interpolation errors above steep orography, perturbations from the reference state (Zängl et al., 2015) rather than the full values are interpolated for the thermodynamic variables ρ and θ_v . In addition, the interpolation operator $\mathcal{I}_{p \rightarrow c}$ is also applied to the model orography in the boundary interpolation zone in the setup phase of the model (before calculating the vertical grid) in order to ensure consistency of the model grids even if different raw data sets have been used to generate the model orographies. To allow a smooth transition to the orography in the interior of the nested domain, a linear blending with a width of eight cell rows (as for the abovementioned nudging zone) is applied.

Regarding the interpolation of edge-based variables (i.e., the edge-normal vector components v_n and ρv_n), we distinguish between *outer child edges* coinciding with the edges of the parent cell and *inner child edges* (see Fig. 2a).

Edge-normal vector components at the inner child edges are reconstructed using a direct RBF reconstruction based upon the five-point stencil indicated by solid red dots in Fig. 2a. For a given inner child edge the stencil comprises the edges of the corresponding parent cell and the two edges of the neighboring parent cells that (approximately) share the orientation of the inner child edge.

For the outer child edges, a more sophisticated reconstruction is applied in order to ensure that the mass flux across a parent edge equals the sum of the mass fluxes across the corresponding child edges. We start with an RBF reconstruction of the 2D vector of the respective variable at the parent triangle vertices (blue triangles in Fig. 2b) using the six (five at

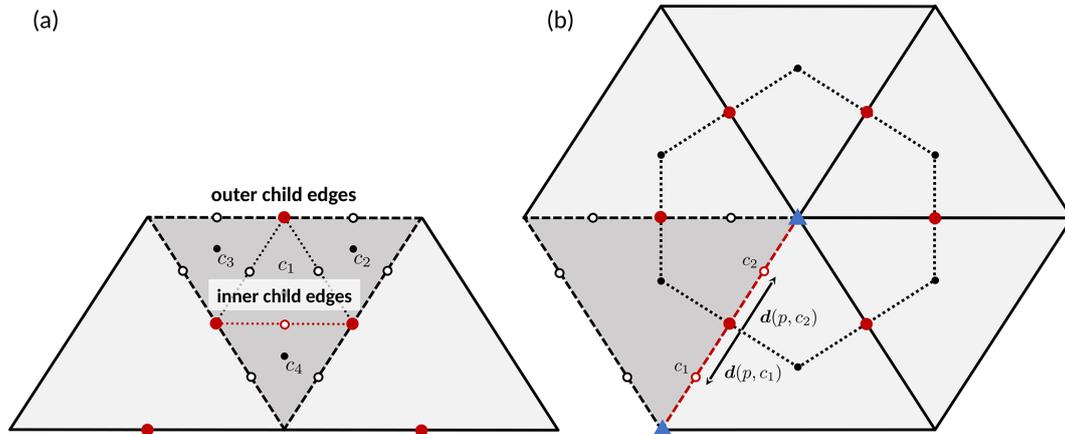


Figure 2. Horizontal reconstruction stencil for edge-normal vector components at (a) inner child edges and (b) outer child edges. The child edge under consideration is highlighted in red. Black open dots indicate child edge midpoints, while black solid dots indicate cell circumcenters. Solid red dots represent the reconstruction stencil, i.e., the location of the parent edge-normal vector components entering the reconstruction, and blue triangles in (b) indicate the location of the reconstructed 2D vectors. See the text for details.

the original vertices of the icosahedron, the *pentagon points*) edge points adjacent to a vertex (red dots).

The edge-normal vector component ϕ at the child edge is then computed as

$$\phi_{c_e} = \phi_p + \nabla_t \phi_p \cdot \mathbf{d}(p, c_e), \quad \forall e \in \{1, 2\}, \quad (2)$$

with $\mathbf{d}(p, c_e)$ denoting the distance vector between the parent and child edge midpoints for a given parent edge and $\nabla_t \phi_p$ denoting the gradient of the edge-normal vector component ϕ_p tangential to the parent edge. The latter is computed by projecting the reconstructed 2D vectors at the two vertices of an edge onto the edge-normal direction and taking the centered difference. Since by construction $\mathbf{d}(p, c_1) = -\mathbf{d}(p, c_2)$ holds on the ICON grid, the abovementioned mass flux consistency is ensured. It is noted that attempts to use higher-order polynomial interpolation methods, which are the standard in mesoscale models with regular quadrilateral grids, were unsuccessful on the triangular ICON grid because the ensuing equation system leads to the inversion of nearly singular matrices.

Rather than interpolating v_n and its time tendency, only the time tendency is interpolated and then used to update v_n at child level at every dynamics time step. The wind field v_n itself is interpolated only once during the initialization of the child domain. This methodology has been chosen because the comparatively inaccurate interpolation to the interior child edges tends to induce small-scale noise in v_n . To suppress the remaining noise arising from the interpolation of the time tendency, a second-order diffusion operator is applied in the inner half of the boundary interpolation zone on v_n , and the default fourth-order diffusion applied in the prognostic part of the model domain (Zängl et al., 2015) is enhanced in up to five grid rows adjacent to the interpolation zone by an amount exponentially decaying with distance

from the boundary. For the second-order diffusion, a coefficient of $0.005 a_e / \Delta t$ is applied, where a_e is the area represented by the current triangle edge, and the scaling with Δt means that the amount of diffusion is effectively independent from the time step. This proved to be sufficient to suppress the development of spurious disturbances even in the data assimilation cycle, which turned out to be the most critical application mode in this respect. For the other prognostic variables, no special filtering is applied near nest boundaries. In the case of one-way nesting, the second-order velocity diffusion is extended into the nudging zone of the nested domain, replacing the enhanced fourth-order diffusion. More details on the nudging zone are given in Sect. 2.2.3.

For the horizontal mass flux ρv_n , the time average over the dynamic sub-steps is interpolated instead of time level n , as the time-averaged mass flux is used by the tracer transport scheme in order to achieve consistency with continuity. Using the mass flux time tendency that is interpolated as well, the related time shift is corrected for when applying the boundary mass fluxes at the child level. In the nested domain, the interpolated mass fluxes valid for the current time step are then prescribed at the interface edges separating the boundary interpolation zone from the prognostic part of the nested domain (edges nr. 9 in Fig. 1). Due to the flux-form scheme used for solving the continuity equation (Zängl et al., 2015), this implies that the interpolated values of ρ do not enter into any prognostic computations in the dynamical core. They are needed, however, for flux limiter computations in the transport scheme. Moreover, no mass fluxes at interior child edges are used so that the nonconservative interpolation method used for those edges does not affect the model's conservation properties. For θ_v and the tracer variables q_k , the values at the edges are reconstructed in the usual manner following Eq. (20) in Zängl et al. (2015) and then multiplied

with the interpolated mass fluxes before computing the flux divergences.

2.2.2 Feedback: child → parent

If two-way nesting is activated, the model state \mathcal{M}_p^{n+1} on the parent domain is relaxed towards the updated model state \mathcal{M}_c^{n+1} on the child domain at every parent fast physics time step Δt_p . This relaxation-type feedback is only applied to the prognostic variables v_n , w , θ_v and ρ plus specific humidity q_v and the specific contents of cloud water q_c and cloud ice q_i . Precipitating hydrometeors are excluded because recommended relaxation timescales (see below) are larger than their typical falling times. Surface variables are excluded as well because they can easily adjust during runtime, and the tile approach used in ICON's land-surface module would require a rather complicated algorithm to avoid inconsistencies.

Let ψ denote any of the above variables. Conceptually, the feedback mechanism is based on the following three steps.

1. *Upscaling.* The updated variable ψ_c^{n+1} is interpolated (upscaled) from the child domain to the parent domain. The upscaling operators for cell-based and edge-based variables will be denoted by $\mathcal{I}_{c \rightarrow p}$ and $\mathcal{I}_{c \rightarrow p}^e$, respectively.
2. *Difference computation.* The difference between the parent domain variable ψ_p^{n+1} and the upscaled child variable $\mathcal{I}_{c \rightarrow p}(\psi_c^{n+1})$ is computed.
3. *Relaxation.* The variable on the parent domain is relaxed towards the upscaled child domain variable by an increment that is proportional to the difference computed in step 2.

For edge-based normal velocity v_n , the arithmetic average of the two child edges lying on the parent edge is taken.

$$\mathcal{I}_{c \rightarrow p}^e(v_{n,c}) = \frac{1}{2} [v_{n,e_{\text{child } 1}} + v_{n,e_{\text{child } 2}}] \quad (3)$$

For cell-based variables the upscaling consists of a modified barycentric interpolation from the four child cells to the corresponding parent cell:

$$\mathcal{I}_{c \rightarrow p}(\psi_c) = \sum_{j=1}^4 \alpha_j \psi_{c_j}. \quad (4)$$

The weights α_j are derived from the following constraints in Eqs. (5)–(7). First of all, a necessary property for the interpolation operator is that it reproduces constant fields, i.e., the weights are normalized:

$$\sum_{j=1}^4 \alpha_j = 1. \quad (5)$$

Moreover, the interpolation shall be linear: with the four child cell circumcenters \mathbf{x}_j ($j = 1, \dots, 4$) and \mathbf{x}_p denoting

the parent cell center, i.e., the interpolation target, we set

$$\sum_{j=1}^4 \alpha_j (\mathbf{x}_j - \mathbf{x}_p) = 0. \quad (6)$$

To motivate this constraint, consider the special case of equilateral triangles in which the center point of the inner child cell \mathbf{x}_1 coincides with the parent center such that the term $(\mathbf{x}_1 - \mathbf{x}_p)$ vanishes. Equation (6) then defines a barycentric interpolation within the triangle spanned by the mass points of the three outer child cells $\{c_2, c_3, c_4\}$ (see Fig. 2a), where the weights $\{\alpha_2, \alpha_3, \alpha_4\}$ represent the barycentric coordinates. Generally, the constraints in Eqs. (5) and (6) ensure reversibility in the sense that $\mathcal{I}_{c \rightarrow p}(\mathcal{I}_{p \rightarrow c})$ returns the original parent cell value irrespective of the reconstructed gradient.

Of course, the contribution of the point \mathbf{x}_1 closest to the interpolation target is of particular importance. Therefore, the underdetermined system of Eqs. (5) and (6) is closed with a final constraint which reads

$$\alpha_1 = \frac{a_{c_1}}{a_p}, \quad (7)$$

where a_{c_1} and a_p denote the inner child and parent cell areas, respectively. In other words, the inner child cell c_1 containing the parent cell circumcenter is given a predefined weight corresponding to its fractional area coverage. This can be interpreted as a conservation constraint for the special case of a very localized signal at the mass point of the inner child cell.

In summary, this method can be regarded as a *modified* barycentric interpolation for the mass points $\{\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4\}$, which accounts for \mathbf{x}_1 as an additional fourth source point. A more stringent barycentric interpolation would require an additional triangulation based on the child mass points.

We note that the cell-based operator $\mathcal{I}_{c \rightarrow p}$ is not strictly mass-conserving and that strict mass conservation would require some means of area-weighted aggregation from the child cells to the parent cells, which is available as an option. The difficulty with such methods on the ICON grid is related to the fact that the mass points lie in the circumcenter rather than the barycenter of the triangular cells. These points coincide on planar equilateral triangles, but they do not on general spherical triangles, the differences being largest in the vicinity of the pentagon points. Due to this fact, using an area-weighted aggregation from the child cells to the parent cells would map linear horizontal gradients on the child grid into a checkerboard noise pattern between upward- and downward-oriented triangles on the parent grid.

Another difficulty that was encountered in the context of mass conservation is related to the fact that the density decreases roughly exponentially with height. In the presence of orography, the atmospheric mass resolved on the model grid therefore increases with decreasing mesh size, assuming the usual area-weighted aggregation of the orographic raw data to the model grid. Feeding back ρ is thus intrinsically nonconservative. To keep the related errors small and non-systematic and to generally reduce the numerical errors over

steep mountains, perturbations from the reference state are used for upscaling ρ and θ_v to the parent grid. A closer investigation of the related conservation errors revealed that the previously mentioned differences between bilinear and area-weighted averaging are negligible (with real orography) compared to the mesh-size-related conservation error.

When combining the abovementioned steps, the feedback mechanism for ρ can be cast into the following form:

$$\rho_p^* = \rho_p^{n+1} + \frac{\Delta t_p}{\tau_{fb}} \left(\mathcal{I}_{c \rightarrow p}(\rho_c^{n+1} - \Delta \rho_{corr}) - \rho_p^{n+1} \right). \quad (8)$$

Here ρ_p^{n+1} denotes the parent cell density, which has already been updated by dynamics and physics. The superscript * signifies the final solution including the feedback increment. Δt_p is the fast physics time step on the parent domain, and τ_{fb} is a user-defined relaxation timescale that has a default value of $\tau_{fb} = 10800$ s and is independent of the relaxed field. The smaller the value of τ_{fb} , the faster the parent state is drawn towards the child state. The chosen default value is optimized for our typical NWP applications and aims at filtering small-scale transient features from the feedback, while fully capturing synoptic-scale features.

Finally note that the upscaled density includes the correction term $\Delta \rho_{corr}$, which has been introduced to account for height differences between the child and parent cell circumcenters. At locations with mountainous orography, the heights at parent cell circumcenters can differ markedly from those at the corresponding child cells. Without an appropriate correction, the feedback process would introduce a noticeable bias in the parent domain's mass field. The correction term is given by

$$\Delta \rho_{corr} = \left(1.05 - 0.005 \mathcal{I}_{c \rightarrow p}(\theta_{v,c}^{n+1}) \right) \Delta \rho_{ref,p},$$

with the parent–child difference of the reference density field

$$\Delta \rho_{ref,p} = \mathcal{I}_{c \rightarrow p}(\rho_{ref,c}) - \rho_{ref,p}$$

and the potential temperature perturbation $\theta_{v,c}^{n+1} = \theta_{v,c}^{n+1} - \theta_{v,ref,c}$. The term $\Delta \rho_{ref,p}$ is a pure function of the parent–child height difference and can be viewed as a leading-order correction term. As a further optimization, the empirical factor $(1.05 - 0.005 \mathcal{I}_{c \rightarrow p}(\theta_{v,c}^{n+1}))$ was added, which means that the correction is roughly proportional to the actual air density. We further note that a possibly more accurate and less ad hoc approach would require a conservative remapping step in the vertical, prior to the horizontal upscaling.

Care is required in order to achieve consistency with continuity. For this purpose, feedback is not implemented for the tracer mass fractions directly, but for partial densities. Building upon the implementation for ρ , we get

$$(\rho q_k)_p^* = (\rho q_k)_p^{n+1} + \frac{\Delta t_p}{\tau_{fb}} \left[\mathcal{I}_{c \rightarrow p} \left((\rho_c^{n+1} - \Delta \rho_{corr}) q_{k,c}^{n+1} \right) - (\rho q_k)_p^{n+1} \right]. \quad (9)$$

Mass fractions are re-diagnosed thereafter:

$$q_{k,p} = \frac{(\rho q_k)_p^*}{\rho_p^*}.$$

By summing Eq. (9) over all partial densities, we recover Eq. (8) for the total density.

A very similar approach is used for θ_v . As for ρ , only the increment of θ_v is upscaled from the child domain to the parent domain and added to the parent reference profile $\theta_{v,ref,p}$.

$$\theta_{v,p}^* = \theta_{v,p}^{n+1} + \frac{\Delta t_p}{\tau_{fb}} \left(\mathcal{I}_{c \rightarrow p}(\theta_{v,c}^{n+1}) + \theta_{v,ref,p} - \theta_{v,p}^{n+1} \right)$$

The same approach is taken for w , but the full field is up-scaled.

$$w_p^* = w_p^{n+1} + \frac{\Delta t_p}{\tau_{fb}} \left(\mathcal{I}_{c \rightarrow p}(w_c^{n+1}) - w_p^{n+1} \right)$$

In the case of v_n some second-order diffusion is added to the ensuing feedback increment in order to damp small-scale noise:

$$v_{n,p}^* = v_{n,p}^{n+1} + \frac{\Delta t_p}{\tau_{fb}} \left(\Delta v_{n,p} + K \nabla^2 (\Delta v_{n,p}) \right), \quad (10)$$

with the feedback increment

$$\Delta v_{n,p} = \mathcal{I}_{c \rightarrow p}^e(v_{n,c}^{n+1}) - v_{n,p}^{n+1}$$

and the diffusion coefficient $K = \frac{1}{12} \frac{a_{p,e}}{\Delta t_p}$, where $a_{p,e}$ is the area of the quadrilateral spanned by the vertices and cell centers adjacent to the parent's edge.

2.2.3 Lateral nudging

If the feedback is turned off, i.e., if one-way nesting is chosen, a nudging of the prognostic child grid variables towards the corresponding parent grid values is needed near the lateral nest boundaries in order to accommodate possible inconsistencies between the two grids, particularly near the outflow boundary. Nudging is performed every fast physics time step of the child domain Δt_c following Davies (1976), by adding a forcing term to the prognostic equations for v_n , ρ , θ_v and q_v of the form

$$\frac{\partial \psi_c}{\partial t} = \text{RHS} + \frac{1}{\Delta t_c} \alpha_{nudge} \left[\mathcal{I}_{p \rightarrow c}(\psi_p - \mathcal{I}_{c \rightarrow p}(\psi_c)) \right],$$

with the nudging coefficient α_{nudge} and the term in brackets denoting the nudging increment. Because lateral boundaries are in general not straight lines on the unstructured ICON grid, attempts to make an explicit distinction between inflow and outflow boundaries (e.g., by prescribing v_n at inflow boundaries only) were not successful.

To compute the nudging increment, the child grid variables are first upscaled to the parent grid in the same way as for the

feedback (Eqs. 3 and 4), followed by taking the differences between the parent grid variables and the upscaled child grid variables. The differences are then interpolated back to the child grid using the same methods as for the lateral boundary conditions (Eqs. 1 and 2). The nudging coefficient α_{nudge} decreases exponentially from the inner margin of the boundary interpolation zone towards the interior of the model domain and is defined as

$$\alpha_{\text{nudge}} = \begin{cases} A_0 \exp\left(-\frac{r-r_0}{\mu}\right), & \text{if } r > r_0 \text{ and } r - r_0 \leq L \\ 0 & \text{otherwise,} \end{cases}$$

with the maximum nudging coefficient A_0 , the cell row index r , the nudging zone start index $r_0 = 5$ (see Fig. 1), the nudging zone width L and the e -folding width μ ; the latter two are defined in units of cell rows. The coefficients A_0 , L and μ may be adjusted by the user, with the default values given by $A_0 = 0.1$, $L = 8$ cell rows and $\mu = 2$ cell rows. A second-order diffusion on v_n is used near the lateral nest boundaries in order to suppress small-scale noise. As opposed to Eq. (10) for feedback, the diffusion operator acts on the velocity field itself rather than the velocity increment.

2.3 Vertical nesting

The vertical nesting option allows setting model top heights individually for each domain, with the constraints that the child domain height is lower than or at most equal to the parent domain height and that the child domain extends into heights at which the coordinate surfaces are flat. This allows, for instance, a global domain extending into the mesosphere to be combined with a child domain that extends only up to the lower stratosphere, which can save significant computational resources. However, a vertical refinement in the sense that the vertical resolution in the child domain may differ from that in the parent domain is not implemented.

Vertical nesting requires appropriate boundary conditions for all prognostic variables to be specified at the vertical nest interface level, i.e., the uppermost half level of the nested domain. This is crucial in order to prevent vertically propagating sound and gravity waves from being spuriously reflected at the nest interface. In the following, boundary conditions are derived for v_n , w , θ_v , ρ and q_k as well as the vertical mass flux ρw . We note that the boundary condition for w is required for the vertically implicit sound wave solver in the dynamics, whereas ρw is needed to compute the vertical flux divergence terms in the prognostic equations for ρ , π and ρq_k .

Due to the constraints mentioned above, boundary conditions can be derived by horizontal parent-to-child interpolation, without the need for any boundary interpolation zone extending vertically away from the upper nest boundary. For w , θ_v , ρ and ρw the full fields at the nest interface level are horizontally interpolated from the parent to the child grid using the same RBF-based interpolation method as for the

lateral boundary conditions (Eq. 1). Rather than interpolating instantaneous values as for the lateral boundaries, w , θ_v , ρ and ρw are averaged over all dynamics sub-steps constituting a fast physics time step in order to filter oscillations related to vertically propagating sound waves. Hence, for $\psi \in \{w, \theta_v, \rho, \rho w\}$ the parent-to-child interpolated field reads

$$\bar{\psi}_c = \mathcal{I}_{p \rightarrow c} \left(\frac{1}{n_{\text{subs}}} \sum_{s=1}^{n_{\text{subs}}} \psi_p^{n+s/n_{\text{subs}}} \right), \quad (11)$$

with s denoting an individual dynamics sub-step and n_{subs} denoting the total number of sub-steps. We apply the same interpolation method to the corresponding time tendencies $\partial \psi_p / \partial t$, which are estimated by taking the difference of the state variables at the sub-steps $s = 1$ and $s = n_{\text{subs}}$. This enables us to perform a linear interpolation in time in order to provide the boundary conditions at approximately the right time levels for every dynamics sub-step on the child domain.

A slightly different approach is taken for v_n , which turned out to be beneficial in order to reduce the magnitude of the horizontal interpolation errors. The differences between the nest interface level and the next half-level below (denoted as $\Delta v_{n,p}$ in the following) are interpolated rather than the full field, again using the same methods as for the lateral boundary conditions. After interpolating $\Delta v_{n,p}$ to the child domain, it is added to $v_{n,c}$ at the second interface level ($k = 3/2$) on the child domain in order to obtain the upper boundary condition, i.e.,

$$v_{n,c}(k = 1/2) = v_{n,c}(k = 3/2) + \mathcal{I}_{p \rightarrow c} \left(\frac{1}{2} \left(\Delta v_{n,p}^n + \Delta v_{n,p}^{n+1} \right) \right).$$

Since Δv_n is less strongly affected by sound waves, only an average between the first and last dynamics sub-step is taken prior to the interpolation. The temporal interpolation is neglected.

For the tracer variables we refrain from interpolating the partial mass fluxes $(\rho w q_k)_p$ directly in order to ensure tracer and air mass consistency. Instead, we make use of the vertical mass flux boundary condition $(\overline{\rho w})_c$ and multiply it with proper mass fractions. On the parent domain the required mass fractions are derived by taking the ratio of the vertical tracer mass flux at the nest interface level calculated in the vertical tracer transport scheme $(\rho w q_k)_p$ and the available mass flux $(\rho w)_p$. The mass fractions are then interpolated to the child domain using Eq. (1). Hence, the flux boundary condition for an arbitrary tracer field q_k reads

$$(\rho w q_k)_c = (\overline{\rho w})_c \mathcal{I}_{p \rightarrow c} \left(\frac{(\rho w q_k)_p}{(\rho w)_p} \right),$$

with $(\overline{\rho w})_c$ computed via Eq. (11.) We note that due to the lack of q_k values above the nest upper boundary (lack of a boundary interpolation zone), the flux computation for scalars at the second interface level ($k = 3/2$) is only stable for vertical Courant numbers $C_k = |w_k| \Delta t_c / \Delta z_k|_{k=3/2} \leq 1$, with Δz denoting the vertical layer thickness.

2.4 Recursive algorithm for multi-domain setups

The previous sections have focused on the coupling of a single child domain. The nesting capability of ICON, however, is not limited to a single domain but supports multiple nests at the same level and multi-level nesting, as well as a combination of both. In the literature, multi-level nesting is also referred to as telescoping nesting (Mouallem et al., 2022). An example of multiple same-level nesting will be provided in Sect. 3.1, while for multi-level nesting we refer to the ICON simulations in Weimer et al. (2021), wherein a three-domain three-level setup has been used to investigate mountain-wave-induced polar stratospheric clouds.

The coupling of multiple same-level nested domains with a parent domain is rather straightforward, as it only requires the single-nest coupling strategy (Sect. 2.2.1–2.2.2) to be applied sequentially for each nest. The coupling strategy for repeatedly nested domains is probably less obvious and will be described here for clarity and in order to complement already existing applications of this feature (e.g., Weimer et al., 2021).

Figure 3 displays a basic multi-level nesting example, wherein a global domain is combined with two successively (two-way) nested domains. The global domain is indicated at the bottom, and the nested domains are vertically staggered on top of it. The red and blue regions show the boundary interpolation zones and feedback zones of the individual domains, respectively. The integration time step on the global domain is denoted by Δt . It is automatically reduced by a factor of 2 when progressing to the next child grid level.

The whole processing sequence for integrating all domains from time step n to $n+1$ is shown in the flowchart at the lower left of Fig. 3. The domains are ordered top-down. Open and filled black dots show model states without and with feedback increments included, black arrows indicate time integration with the fast physics time step, and red and blue arrows indicate lateral boundary data interpolation and feedback, respectively.

The flow control of ICON's hierarchical nesting scheme is handled by a recursive subroutine that cascades from the global domain (or outer limited-area domain) down to the deepest nesting level and calls the dynamical core and the physics parameterizations for each domain in basically the same way as for the global domain. The basic processing sequence is as follows.

1. A single integration step with Δt is executed on the global domain, resulting in an updated model state \mathcal{M}_p^{n+1} , as indicated by an open black circle in Fig. 3.
2. Boundary data are interpolated from the global domain to nest 1 (red arrow), followed by an integration step on nest 1 over the time interval $\Delta t/2$ and resulting in the model state $\mathcal{M}_{c1}^{n+1/2}$.

3. As another nested domain exists within nest 1, boundary fields based on the model state $\mathcal{M}_{c1}^{n+1/2}$ are interpolated to the second nested domain. Afterwards, the model is integrated on nest 2 over 2 times the time interval $\Delta t/4$, resulting in the model state $\mathcal{M}_{c2}^{n+1/2}$.
4. Feedback is conducted from nest 2 back to nest 1 (blue arrow), resulting in an updated model state $\mathcal{M}_{c1}^{n+1/2*}$ on nested domain 1 (black filled dot). Then, on the nested domain 1 the model is again integrated in time to reach model state \mathcal{M}_{c1}^{n+1} .
5. This is followed by a second lateral boundary data interpolation from nest 1 to nest 2 based on \mathcal{M}_{c1}^{n+1} . Nest 2 is integrated in time again to reach its state \mathcal{M}_{c2}^{n+1} .
6. As a final step, feedback is performed from nest 2 to nest 1, followed by feedback from nest 1 to the global domain.

We note that the presented coupling strategy is very similar to that in WRF or FV3, but differences exist in several details. For example, the child-to-parent feedback in FV3 covers only temperature and the wind components (Mouallem et al., 2022), which avoids any impact on mass conservation but apparently necessitates executing the feedback before the physics call at the parent level in order to maintain numerical stability.

3 Application examples

3.1 Jablonowski–Williamson test

To demonstrate the functionality of the grid nesting and to investigate the numerical errors related to the mesh size discontinuity along the nest boundary, we start with the Jablonowski and Williamson (2006) test (called the JW test hereafter) already used in Zängl et al. (2015) and many other studies to evaluate basic aspects of numerical accuracy. This test considers the formation of baroclinic waves in a geostrophically and hydrostatically balanced zonally symmetric basic state with a very strong Equator-to-pole temperature contrast. In its standard version, the waves are triggered by a small perturbation in the initial wind field imposed at 20° N, 40° E. The perturbation grows very slowly during the first few days but evolves into an explosive cyclogenesis between days 7 and 10 of the test case. Dropping the initial perturbation, which implies that the model ideally should maintain the initial state, allows for testing the numerical errors related, for instance, to grid irregularities.

Our first series of experiments considers the baroclinic wave test for a variety of configurations based on the model grids R2B4 (mesh size 160 km) and R2B5 (80 km). Experiment E1 uses the global R2B4 grid only, experiments E2–E4 use a global R2B4 grid plus nested R2B5 grids at various locations along the track of the baroclinic wave (see Table 1 for

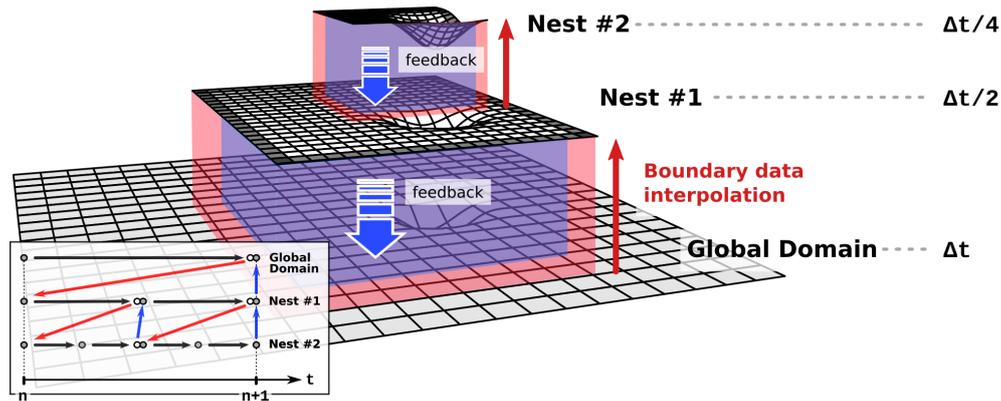


Figure 3. Schematic of a multi-domain multi-level setup with two domains nested successively into a global (or limited-area) base domain. The processing sequence for the time integration of all domains from time step n to $n + 1$ is shown in the flowchart at the lower left. See the text for details.

Table 1. List of model configurations for the Jablonowski–Williamson test.

Exp. ID	Global grid	Position of first nest	Position of second nest
E1	R2B4	–	–
E2	R2B4	32.5–72.5° N, 65–115° E	–
E3	R2B4	32.5–72.5° N, 65–115° E	32.5–72.5° N, 175–125° W
E4	R2B4	32.5–72.5° N, 30° E–130° W	–
E5	R2B5	–	–

details), and E5 uses a global R2B5 grid for reference. To account for the fact that these resolutions are much coarser than in typical NWP applications, the relaxation timescale τ_{fb} is increased to 12 h. The results are summarized in Fig. 4, showing the surface pressure and the relative vorticity at 850 hPa. For the nested experiments E2–E4, the results are displayed for the coarse (R2B4) grid, implying that the impact of the nested domain(s) appears indirectly via the feedback mechanism described above. Selected results for the nested domain are added in Fig. 5, while combined results are displayed in Fig. 6, where the highest-resolution data are used in any region of the plot.

As discussed in Zängl et al. (2015), the general behavior of ICON (and most other grid point models) in the JW test is that the baroclinic wave train exhibits a phase lag at coarse resolution. Moreover, the leading cyclone tends to be too weak at coarse resolution, whereas the trailing cyclone tends to be slightly too intense because grid irregularities act as an additional trigger for wave-like disturbances. In Fig. 4, these features become evident from comparing E1 (Fig. 4a) with E5 (Fig. 4i). Comparing the nested runs E2 and E3 (Fig. 4c, e) with E1 reveals that the western nested domain, which is shared by E2 and E3 and is crossed by the baroclinic wave during the early stage of its evolution (days 2–4), leads to a slight reduction of the phase lag and the intensity bias of the trailing cyclone. On the other hand, there is very little difference between E2 and E3, implying that the eastern nested

domain entered by the leading cyclone at the beginning of its rapid intensification around day 7 has only a weak impact on the cyclone intensity. More pronounced differences appear with the large nested domain of E4, covering the baroclinic wave train during most of its life cycle. In the nested domain (see Fig. 6), the position and intensity of the cyclones are now very similar to E5. The relaxation-based feedback to the coarse domain shown in Fig. 4g shifts the cyclones to the right positions, but their intensity is weaker than in the reference experiment E5. The intensity difference is primarily related to the abovementioned long relaxation timescale of 12 h, combined with the fact that the middle and left cyclones propagate at a somewhat slower speed in the coarse domain. Reducing τ_{fb} gradually decreases the difference, but for values less than about $5\Delta t_p$ (i.e., 2 h), the solution in the nested domain starts to deviate from the E5 reference because numerical disturbances generated along the nest boundaries become non-negligible (not shown). Note in this context that in typical NWP applications, the resolution in the global domain is always high enough that cyclones and fronts move at the same speed as in the nested domain(s) and possess about the same intensity (except for tropical cyclones), so the intensity bias encountered here is not of practical relevance. On the other hand, the standard relaxation timescale of 3 h filters transient wave-like motions that are not deterministically predictable anyway, thus avoiding a transfer of small-scale noise into the global domain.

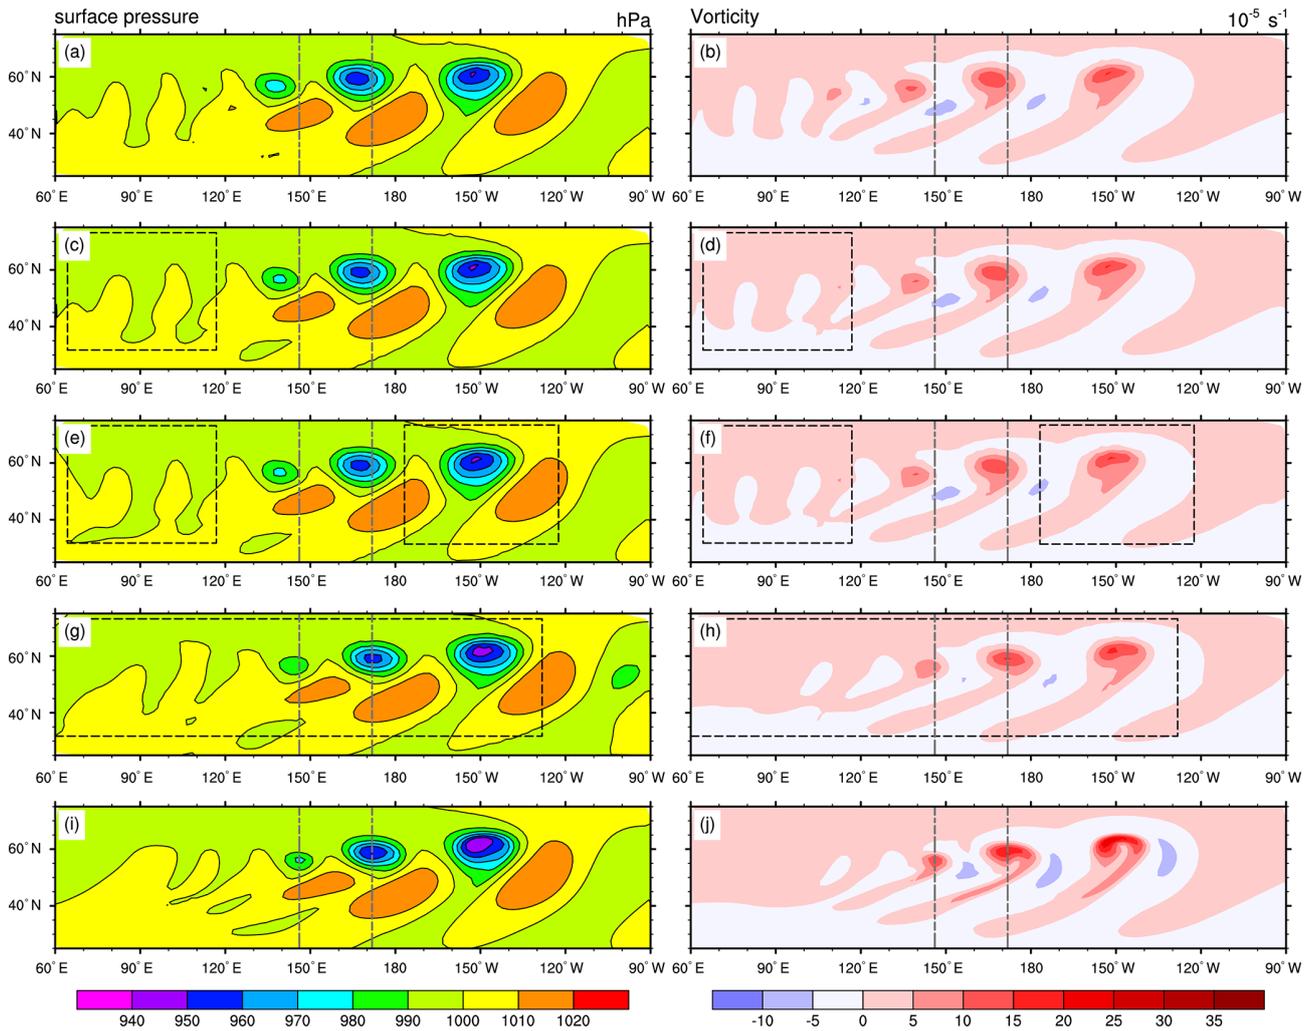


Figure 4. Surface pressure (hPa, left column) and relative vorticity at 850hPa (10^{-5} s^{-1} , right column) for the Jablonowski–Williamson experiments E1 (a, b), E2 (c, d), E3 (e, f), E4 (g, h) and E5 (i, j) after 9 d of integration. Dashed boxes show the nest locations. To ease comparison, the dashed gray line indicates the longitudinal position of the trailing cyclone’s minimum surface pressure in the R2B5 reference run (E5). See Table 1 for the experiment configurations.

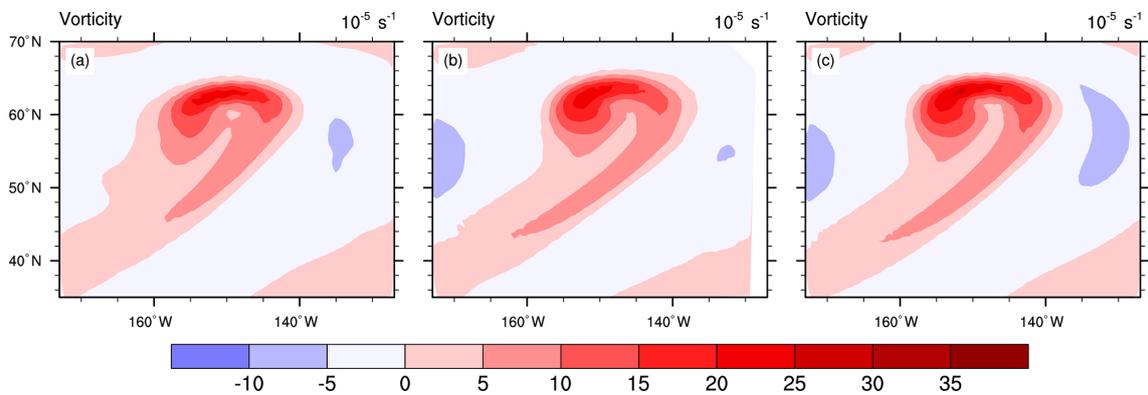


Figure 5. Relative vorticity at 850hPa (10^{-5} s^{-1}) after 9 d of integration for experiments E3 (a), E4 (b) and E5 (c), calculated at R2B5 resolution in each case (i.e., in the nested domain for E3 and E4).

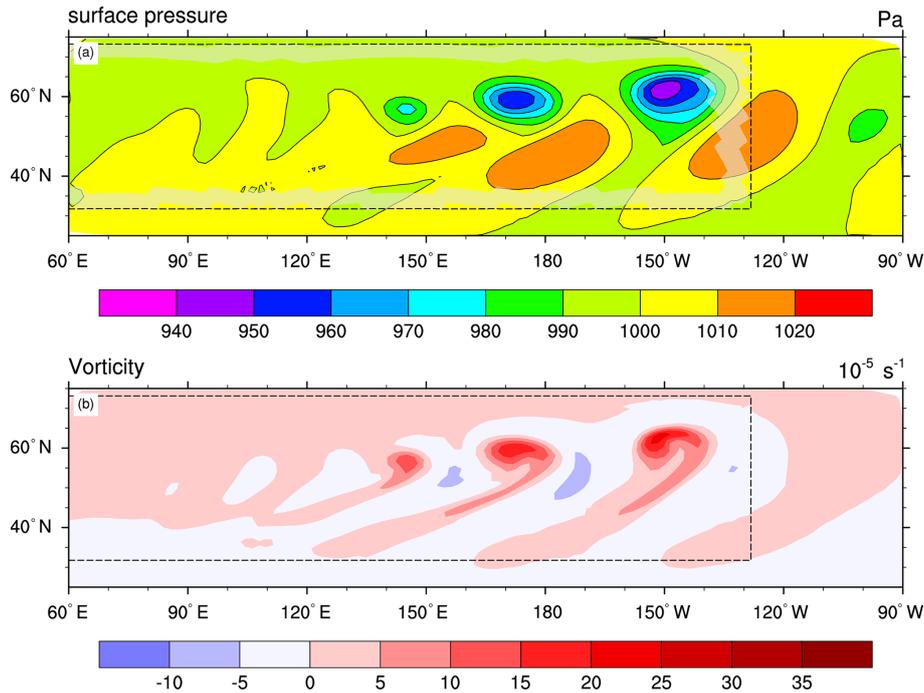


Figure 6. (a) Surface pressure and (b) relative vorticity at 850 hPa (10^{-5} s^{-1}) for JW experiment E4. The combined solution is plotted, which means that the highest-resolution data are used in any region of the plot. The dashed box shows the location of the nest, and the white-shaded region in (a) highlights the exact location of the non-prognostic boundary interpolation zone.

The relative vorticity fields (right panels of Fig. 4) confirm that the nested domain in E2 (and E3) reduces the phase lag of the trailing cyclone as well as the intensity of an extra disturbance on the upwind side that does not appear in E5. However, the intensity differences between E5 and the other experiments appear much more pronounced than in the surface pressure field. This is primarily because the accuracy at which derivative-based quantities like vorticity can be calculated depends directly on the mesh size of the model grid. To obtain better comparability, Fig. 5 displays vorticity fields for the leading cyclone computed at the same mesh size (R2B5) for E3, E4 and E5. Compared to the reference result from E5 (Fig. 5c), the E3 vorticity field exhibits significant distortions and a somewhat reduced amplitude (Fig. 5a), reflecting the fact that the cyclone was only poorly resolved during part of its development stage. For E4 (Fig. 5b), the differences are much smaller and concentrate in the vicinity of the nest boundary at 130° W . Note that there are no evident disturbances near the transition between the lateral boundary interpolation zone and the prognostic computational domain of the nested domain at about 135° W . Further, there is no evidence of significant discontinuities along the parent–child domain interface, as can be deduced from the combined solution plot in Fig. 6.

To further examine the flow disturbances generated by the resolution jump at the nest boundary, Fig. 7 compares the surface pressure fields for the steady-state JW test after 9 d of integration for E1 and E2. This variant of the JW test ac-

centuates the so-called grid imprinting errors arising from any irregularity in the model grid, in our case due to both the nonuniformity of the icosahedral grid and the resolution jump along the nest boundaries. As a side note we mention that previous non-nested ICON results presented in Lauritzen et al. (2010) were obtained with an early version of the hydrostatic dynamical core (Wan et al., 2013), which did not perform as well as the nonhydrostatic dynamical core presented here. The result for E1 (Fig. 7a) shows the well-known regular wavenumber-five disturbance pattern characteristic for icosahedral grids (see, e.g., Jablonowski and Williamson, 2006; Lauritzen et al., 2010). With the presence of a nested domain, the disturbances become irregular and reach a somewhat stronger peak amplitude (Fig. 7b). The largest disturbances occur in the air mass that was initially located in the nest region because the nest-induced perturbations there have the longest time to grow. For smaller values of τ_{fb} , the amplitude of these disturbances gradually increases (not shown). Regarding their practical relevance, we note that the baroclinicity of the real atmosphere is generally much weaker than in the initial state of the JW test because it is continuously depleted by synoptic-scale disturbances. Moreover, the nest-induced perturbations decrease about proportionally to the grid imprinting in the global grid (see Fig. 2 in Zängl et al., 2015) when refining the model resolution. In DWD’s operational NWP applications, we have never encountered noticeable artificial disturbances along nest boundaries.

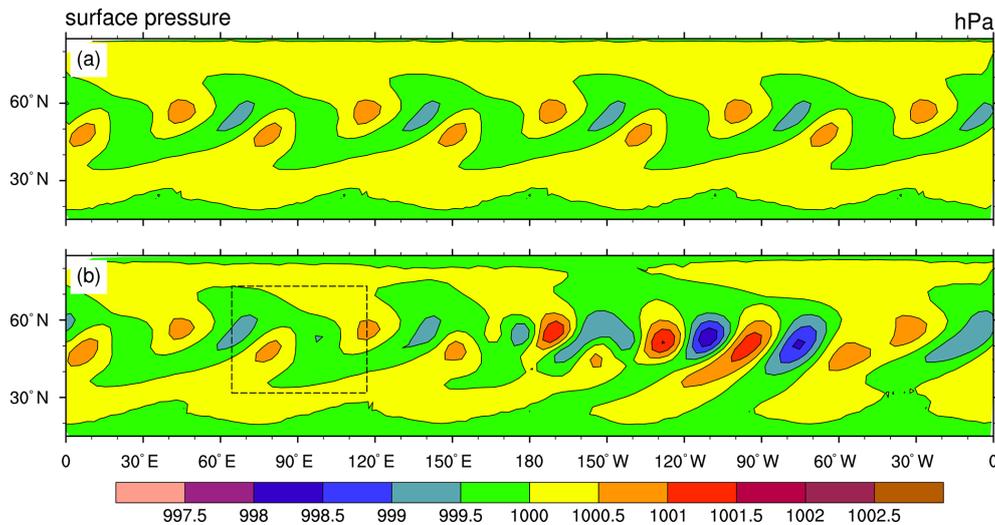


Figure 7. Surface pressure (hPa) after 9 d of integration for the steady-state JW test and experiments E1 (a) and E2 (b). The dashed box shows the nest location.

3.2 Schär mountain test with vertical nesting

In order to examine the behavior of the upper boundary condition for vertically nested domains, we conducted the Schär mountain test case (Schär et al., 2002). Therein, a uniform flow of constant wind speed U and Brunt–Väisälä frequency N over idealized hilly terrain gives rise to a combination of small-scale, vertically decaying nonhydrostatic gravity waves and larger-scale vertically propagating quasi-hydrostatic gravity waves. We expect this test case to challenge the upper boundary condition implementation, as the vertically propagating waves should pass through the domain interface without spurious reflections or any accumulation of noise.

Following Zängl et al. (2015), we performed a quasi-three-dimensional analogue of this test on a limited-area domain, where the original 1D mountain profile is changed to a 2D ridge-like profile that decays towards zero in the cross-flow direction as the lateral domain boundary is approached. In our earlier study, the limited-area setup was chosen in favor of a small planet configuration (Klemp et al., 2015) because it allows for a better control on the upstream flow conditions in the presence of high mountains, and it is retained here for convenience. The ridge-like profile is given by

$$h(x, y) = h_m \exp\left(-\frac{x^2}{a^2}\right) \cos^2\frac{\pi x}{\lambda} \exp\left(-\frac{\max(0, |y| - \beta)^2}{a^2}\right),$$

with the originally proposed parameter settings $h_m = 250$ m, $a = 5000$ m and $\lambda = 4000$ m, as well as the ridge length scale $\beta = 10^5$ m. The wind speed is set to $U = 10$ m s⁻¹, and the Brunt–Väisälä frequency is given by $N = 0.01$ s⁻¹ for $z < 20$ km and gradually increases to $N = 0.03$ s⁻¹ above. Note that the latter setting differs from the constant N used by Schär et al. (2002) in order to allow for the high model top

desired for our test configuration (see below). This does not allow for comparing against an analytic solution, but this is not needed here because model results serving this purpose have already been provided by Zängl et al. (2015).

We have performed two types of simulations: one with and one without a vertically nested domain. The reference simulation without a vertical nest was performed on a single $3^\circ \times 3^\circ$ limited-area domain centered at the Equator with a horizontal mesh size of $\Delta x \approx 620$ m (R2B12), a vertically stretched grid with 50 vertical layers and a 40 km model top. To avoid the reflection of gravity waves at the model top, a Rayleigh damping layer acting on vertical velocity w was applied above $z = 22$ km (i.e., encompassing the uppermost 10 vertical layers). The simulation was run for 12 h until an approximately steady state was reached. The steady-state solution for w is depicted in Fig. 8a. While the wave structure at low levels closely matches previously published results for configurations with constant N (e.g., Schär et al., 2002; Skamarock et al., 2012; Zängl et al., 2015) partial wave reflections probably related to the stability change starting at $z = 20$ km become apparent in the upper part of the domain.

For the nested simulation, a $2.4^\circ \times 2.4^\circ$ R2B12 domain with 39 vertical layers and a vertical interface at 20 km is two-way nested into a $3^\circ \times 3^\circ$ R2B11 parent domain ($\Delta x \approx 1.2$ km) with 50 vertical layers and a 40 km model top. There is no Rayleigh damping active in the nested domain. Note that the vertical layer distribution and the horizontal mesh size of the nested domain exactly match those of the reference simulation. The only significant difference lies in the much lower but open domain top. Steady-state solutions for the nested domain are depicted in Fig. 8b and c. They differ in terms of the relaxation timescale τ_{fb} for child-to-parent feedback, with values of $\tau_{fb} = 10800$ and $\tau_{fb} = 900$ s for panels (b) and (c), respectively.

In general, the results for the nested domain and the reference result are fairly close to each other. There is no indication of substantial wave energy reflection or noise accumulation along the nest interface level. Deviations from the reference result are largely confined to the uppermost quasi-hydrostatic wave crest and trough and to the leeward propagating wave signal. This holds not only for the steady-state solution, but also for the spin-up phase (not shown). The reason for the deviations is twofold: firstly, the computation of the boundary conditions at the nest interface level inevitably goes along with spatiotemporal interpolation errors. Second, and more importantly, the solutions on the parent and child domains are slightly different, implying that the vertical interface condition derived from the parent domain cannot exactly match the solution on the child domain. These differences primarily originate from the differences in the mesh size, but they additionally depend on the feedback timescale. Reducing the feedback timescale strengthens the domain coupling and reduces the parent–child differences, which in turn improves the vertical nest interface conditions. This can be seen by comparing Fig. 8b ($\tau_{fb} = 10800$ s) with Fig. 8c ($\tau_{fb} = 900$ s), which shows that a shorter feedback timescale further reduces the difference to the non-nested reference (Fig. 8a). The corresponding solution on the R2B11 parent domain (with $\tau_{fb} = 900$ s) is depicted in Fig. 9. Due to its 2 times coarser mesh size it is lacking some of the leeward-propagating wave signals and shows a slightly reduced amplitude of the quasi-hydrostatic wave. The attenuated, though visible, leeward-propagating wave signal in Fig. 9 results from the child-to-parent feedback and does not exist in an R2B11 simulation without a nest (not shown).

From this test it can be concluded that, even without any interpolation error, the boundary conditions at the nest interface will never match perfectly due to the resolution-induced differences between the parent and child model states. On the other hand, this test has shown a small but noticeable positive impact of the child-to-parent feedback mechanism on the quality of the nest interface conditions, which improve with decreasing feedback timescale.

3.3 Operational NWP applications

In an operational context, one ideally expects the beneficial impact on forecast quality of the regionally refined resolution to be transferred to the global domain in the nest overlap area and subsequently propagate downstream, which is usually eastward in the extratropics. This implicitly assumes that the scores used to quantify forecast quality do improve with increasing model resolution, which, according to our experience, is the case for mesh sizes coarser than about 10 km but not necessarily in the convective gray zone. In the operational global forecasting system of DWD, the deterministic part has a global horizontal mesh size of 13 km and a two-way nested domain referred to as “ICON-EU” with 6.5 km mesh size covering Europe and some adjacent regions (24.5° W–

63.5° E, 29–71° N; i.e., domain 2 in Fig. 1). For this resolution range, the impact of the nested domain on the global forecast quality tends to be rather small. A much clearer impact is found for the corresponding ensemble prediction system (EPS), which (at the time of writing this paper) uses mesh sizes of 40 and 20 km. To demonstrate the benefit of two-way nesting on NWP quality, we therefore use the EPS configuration of ICON, but the subsequent experiments are executed as deterministic forecasts for simplicity and easier reproducibility outside the operational environment of DWD. Moreover, our experiments are initialized with interpolated operational IFS analyses available from the European Centre for Medium-Range Weather Forecasts (ECMWF), which has the conceptual advantages that none of our experiments may benefit from being initialized from its “own” assimilation cycle and we can verify our forecast results against independent (i.e., IFS) analyses.

The current operational configuration at DWD uses 90 vertical levels with a model top at 75 km and a vertically nested EU domain with 60 levels and a vertical interface to the global domain at about 23 km. The forecast lead time is 180 h in the global domain and 120 h in the nested one. These settings are also used for our subsequent experiments unless stated otherwise. Specifically, we consider our nested EPS configuration, henceforth denoted as R2B6N7, a global 40 km mesh without a nest (R2B6) and a global 20 km mesh without a nest (R2B7). All experiment suites are conducted for January 2021, starting from the 00:00 UTC IFS analysis for each day of the month. The physics configuration equals the operational status of autumn 2021 as far as possible, notable exceptions being that no ensemble perturbations are used and that some tuning options relying upon coupling with the operational data assimilation scheme are turned off (see Reinert et al., 2021, for further details on the ensemble perturbations). Besides the abovementioned verification against IFS analyses, we use DWD’s operational verification against SYNOP (surface) stations and radiosonde ascents (TEMP), following the WMO standard in all cases (WMO, 2019).

To exemplarily demonstrate the resolution dependence of the forecast quality and the related nest impact, Fig. 10 shows the root mean square error (RMSE) against IFS analyses for 500 hPa geopotential and vector wind in the Northern Hemisphere. It is clearly evident that the errors of R2B7 are smaller than for R2B6 during the whole forecast range. The nested setup R2B6N7 exhibits slightly smaller errors than R2B6 and is, as may be expected due to the size of the refined domain, closer to R2B6 than R2B7. A closer look at Europe is given in Fig. 11, showing relative differences to R2B6 for easier readability and including an additional test in which the nest stays active until 180 h (R2B6N7-180h). It is seen that during the first 3 forecast days, the quality improvement related to the nest feedback in R2B6N7 is comparable to what is obtained with a global 20 km mesh (R2B7). Afterwards, the improvement in R2B6N7 starts to decrease as the

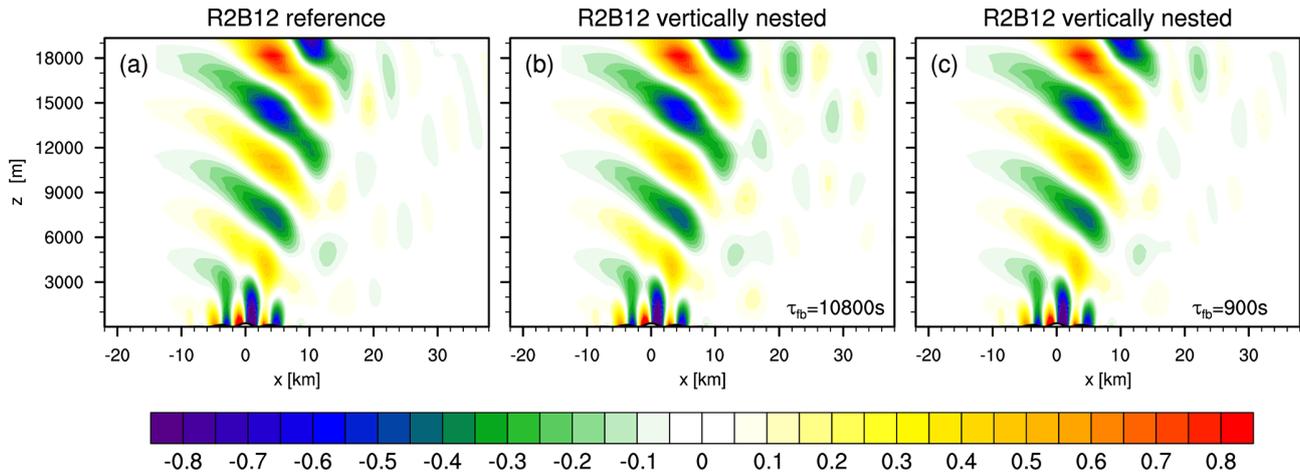


Figure 8. Vertical velocity w after 12h for the Schär et al. (2002) test case (contour interval 0.05 ms^{-1}). (a) Reference results for a single, non-nested R2B12 domain with 40 km model top. (b, c) Results for a vertically nested R2B12 domain with 20 km model top and two distinct relaxation timescales for child-to-parent feedback. The entire vertical extent of the nested domain is depicted.

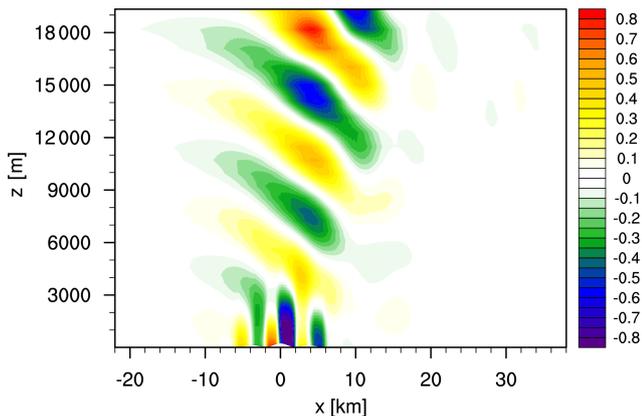


Figure 9. Vertical velocity w after 12h for the R2B11 parent domain, with $\tau_{fb} = 900\text{ s}$.

advection from the coarse (global) domain into the nesting region becomes more and more relevant. After 5 d, the errors in R2B6N7 become temporarily a bit larger than in R2B6 over Europe, and we infer from the negligible differences between R2B6N7 and R2B6N7-180h that this is unrelated to the usual termination of the nest at 120h.

More detailed information on the nest impact is provided by the SYNOP and TEMP verification, which includes an assessment of statistical significance at the 95 % level. The SYNOP verification (Fig. 12) shows a pronounced beneficial impact in the nest overlap region (EU-NEST) during the first 5 forecast days, which is largest for surface pressure (PS) with an improvement of almost 10 % on day 2 and still substantial for 10 m wind speed (FF), 2 m temperature (T2M) and 2 m humidity (RH2M). After the termination of the nested domain, the differences become small and insignificant. When the nest runtime is extended to 180h, the

improvements continue to be present for T2M and RH2M but are negligible for PS and FF (not shown). The results for Asia show that the improvements related to the EU-NEST propagate downstream with a delay of about 3 d. They are a bit smaller than in the nest region but still partly reach the 95 % significance level, indicating an obvious benefit of our two-way nesting methodology. A more detailed view of the spatiotemporal evolution of the forecast skill improvement related to the nest feedback is provided in Fig. 13, showing a station-wise verification of surface pressure against all SYNOP stations used in the operational data assimilation cycle at DWD. On forecast days 1 and 2, the surface pressure RMSE is almost entirely smaller for R2B6N7 than R2B6, exceptions being restricted to a few individual stations. Later on, the spatial variability increases, with the first coherent regions with degradations appearing after about 3 d. Nevertheless, the impact of the nest stays predominantly positive, as seen in Fig. 12, and the eastward propagation of the average skill improvement is confirmed.

The TEMP verification (Fig. 14) additionally shows that the positive signal extends throughout the troposphere in both Europe and Asia. In the lower stratosphere, a slight degradation can be seen over Europe for wind speed and direction as well as temperature, which might be suspected to be related to the vicinity of the vertical nest interface at about 30hPa. However, an additional sensitivity test with the nest extending up to 75 km (not shown) revealed that this is not the case. In this test, the degradation actually extends further up into the middle stratosphere, indicating that it is probably a double-penalty effect related to resolving a larger part of the gravity wave spectrum, an issue we generally observe when increasing the horizontal model resolution (but the discussion of which is beyond the scope of this paper). On average over the Northern Hemisphere, the improvements related to

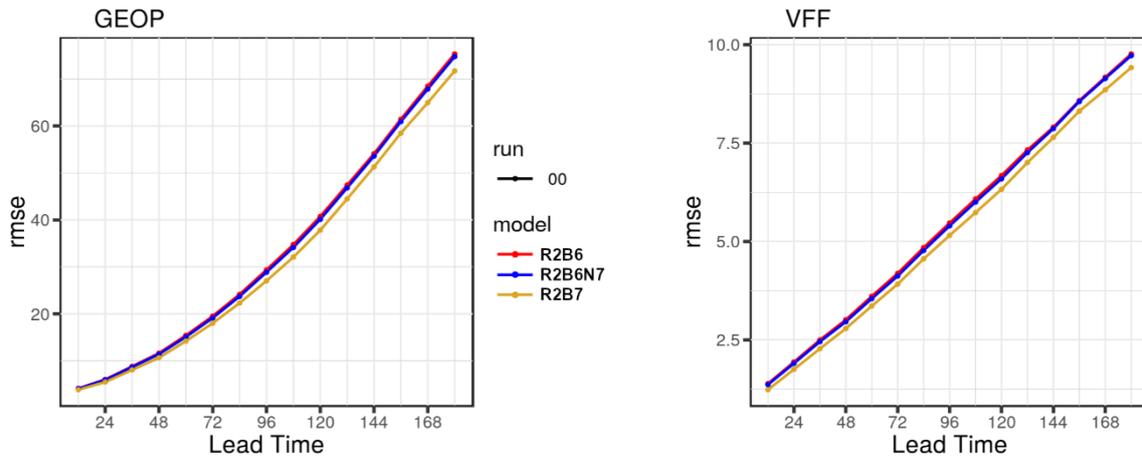


Figure 10. Verification against IFS analyses for the Northern Hemisphere (latitude $> 20^{\circ}$ N) at 500 hPa: RMSEs of geopotential (m) and vector wind (ms^{-1}) as a function of lead time (h). See text for experiment acronyms.

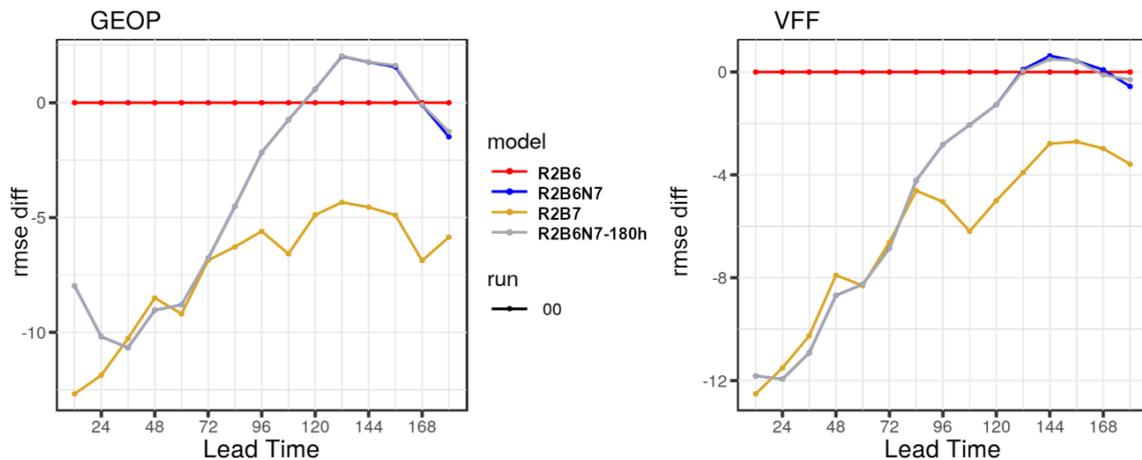


Figure 11. Verification against IFS analyses for Europe ($35\text{--}72^{\circ}$ N, 12° W– 45° E) at 500 hPa: relative rms difference (%) of geopotential and vector wind with respect to R2B6. Note that R2B6N7 and R2B6N7-180h are identical up to a lead time of 120 h. See text for experiment acronyms.

the EU-NEST are still significant for all variables, the magnitude in the TEMP verification being largest for wind speed and direction. We finally mention that the improvement of upper-air geopotential is smaller than for PS, which appears to be a contradiction but is related to the fact that the density of radiosonde stations is more homogeneous than that of the surface stations, giving Europe an overly large weight in the SYNOP verification. Nevertheless, this does not affect our conclusion that the regional refinement over Europe has a clear and significant benefit for the forecast quality in the Northern Hemisphere, demonstrating that the two-way nesting capability in ICON fully serves its purpose.

Further analysis has been undertaken to investigate potential numerical issues related to the nesting, such as possible artifacts along the lateral boundaries of the nested domain and the loss of exact mass conservation mentioned in Sect. 2. Regarding boundary artifacts, precipitation is known

to be one of the most sensitive variables because unphysical flow convergences or divergences tend to induce line-shaped structures along the boundaries that are readily recognized as unrealistic. To consider this aspect, we selected a forecast run in which substantial precipitation amounts occurred along the lateral boundaries (15 January 2021). The result after 5 forecast days is displayed in Fig. 15 for the nested domain and the immediate surroundings. The boundary interpolation zone, in which no prognostic precipitation is calculated in the nested domain, has been filled with interpolated values from the global domain. Although it is visible at a few spots that the interpolated precipitation along the boundaries is smoother than the prognostic one in the interior of the domain (as expected from the coarser model resolution at which it has been calculated), there is no evident discontinuity pointing to the presence of numerical problems.

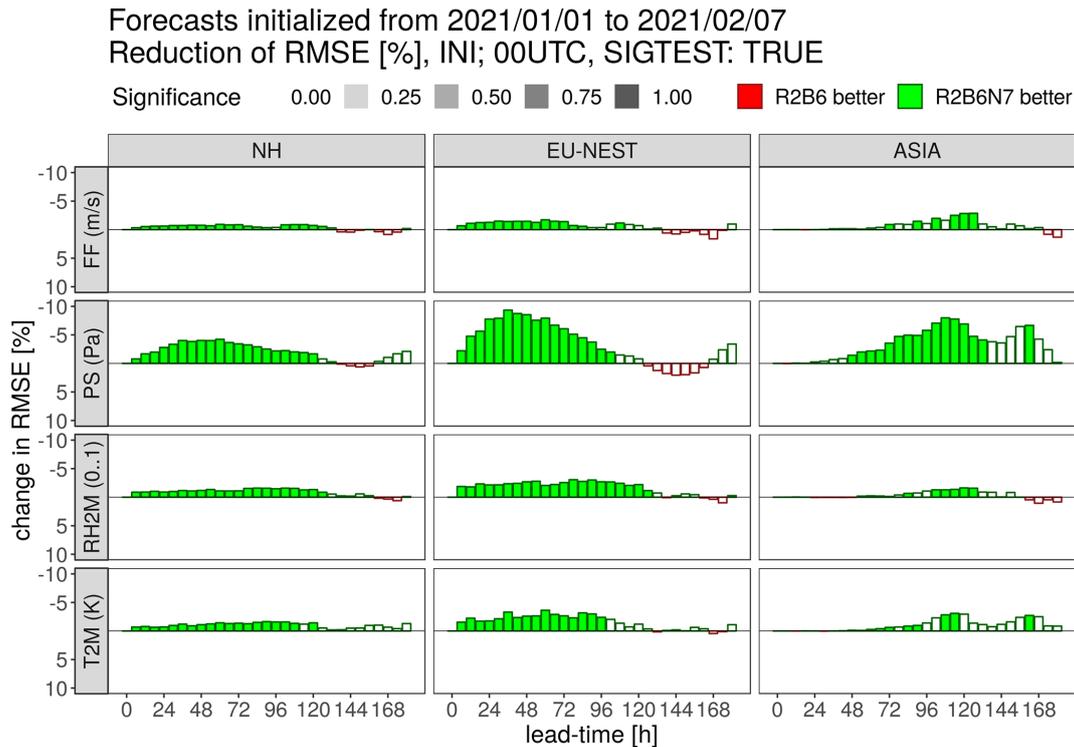


Figure 12. Scorecard for verification against SYNOP observations for 10 m wind speed (FF), surface pressure (PS), 2 m relative humidity (RH2M) and 2 m temperature (T2M) in the Northern Hemisphere (NH), the EU-NEST overlap region (30–70° N, 20° W–60° E) and Asia (30–70° N, 65–140° E). Bars indicate relative RMSE differences between R2B6N7 and R2B6, and color filling indicates statistical significance at the 95 % level.

To examine the impact of the two-way nesting on mass conservation, an additional set of forecast experiments has been conducted in which the lead time was extended to 30 d and the nested domain stays active until the end. Results for 18 arbitrarily selected initial dates are summarized in Fig. 16, showing the relative mass change in the global domain due to the presence of the two-way nest (the accuracy is close to machine precision otherwise). While there appears to be a slight systematic loss of mass during the first forecast days, which appears to be related to spin-up effects, there is no systematic trend visible later on. The relative errors are on the order of 10^{-6} , which is small enough to be irrelevant for NWP purposes, and the absence of a trend indicates that the remaining conservation accuracy would even be sufficient on climate timescales (although this has not been further examined). Additional experiments indicated that the typical mass change during the first 2–3 d is smaller and not systematically negative when starting from operational analyses (not shown). Thus, no further efforts have been undertaken to investigate the apparent spin-up effect with interpolated initial conditions.

4 Conclusions

This article provides a technical description of the grid nesting implementation in the ICOSahedral Nonhydrostatic (ICON) modeling system. The available options comprise one-way and two-way nesting, with one or more domains per nesting level, and vertical nesting in the sense that the upper boundary of a nested domain may be lower than that of its parent domain. In addition, a limited-area mode is available as a by-product of the grid nesting implementation, which differs from one-way nesting only in the way the lateral boundary conditions are provided. The model-internal flow control basically follows the nesting approach known from mesoscale models like MM5 (Grell et al., 1994) or WRF (Skamarock et al., 2019), but the feedback applies a Newtonian relaxation rather than a direct replacement of the prognostic variables in the overlapping part of the parent domain. The relaxation-type feedback was found to produce slightly better forecast quality in NWP applications and obviates the need to adjust the model orography at the parent grid level. The main innovation that was needed compared to the well-established nesting implementations on quadrilateral grids was the development of appropriate interpolation algorithms for the lateral boundary conditions of the nested grids, as the usual higher-order polynomial interpolation leads to the in-

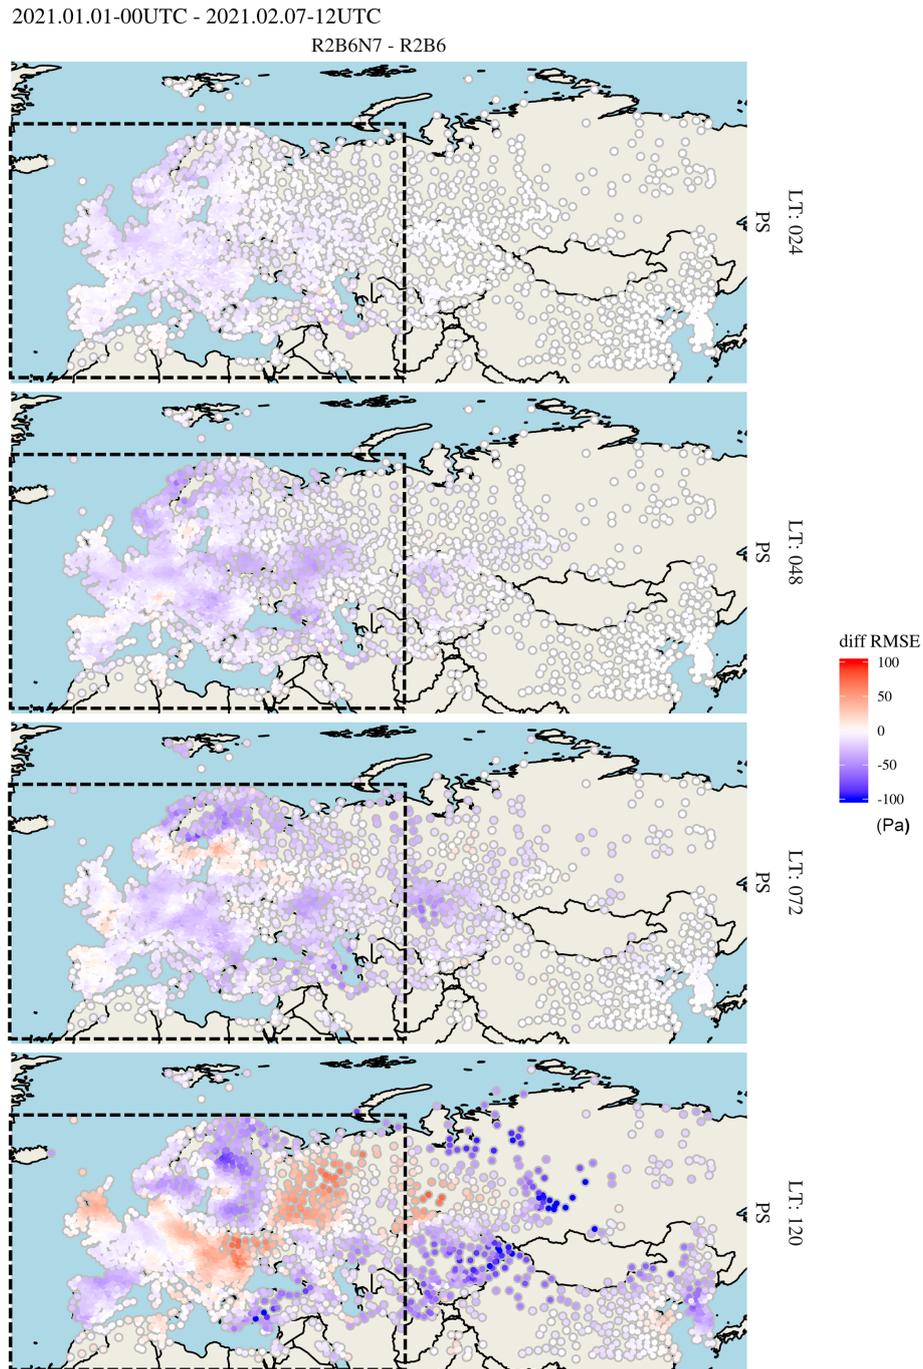


Figure 13. Station-wise verification of the surface pressure RMSE difference (Pa) between experiments R2B6N7 and R2B6 for forecast lead times (LT) of 24, 48, 72 and 120 h. Blue (red) dots indicate smaller (larger) errors for R2B6N7 than R2B6. Dashed boxes show the location of the nested domain.

version of a nearly singular matrix on the triangular ICON grid. It was found that radial basis functions constitute an adequate method. They are used either for reconstructing spatial gradients on the parent grid, followed by a linear extrapolation of the prognostic variables from the parent grid to the

child grid, or for directly reconstructing the variable values on the child grid points.

To demonstrate the functionality and quality of the grid nesting in ICON, idealized tests restricted to the dry dynamical core are presented as are real forecast experiments using a model configuration very close to that used for op-

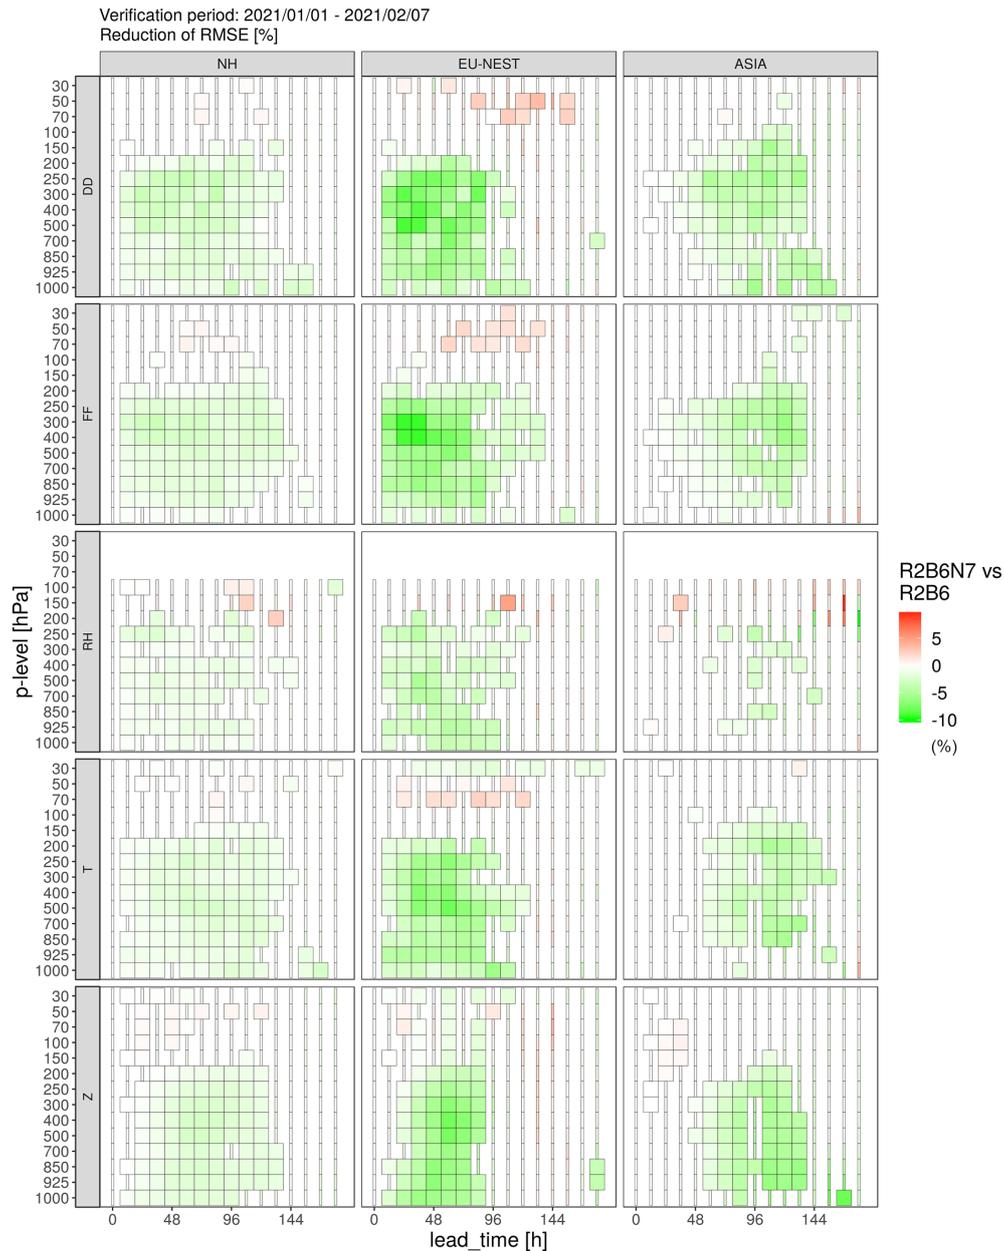


Figure 14. Scorecard for verification against radiosonde ascents for wind direction (DD), wind speed (FF), relative humidity (RH), temperature (T) and geopotential (Z) in the same regions as in Fig. 12. Colors indicate relative RMSE differences (%) between R2B6N7 and R2B6, and wide boxes indicate statistical significance at the 95 % level.

erational NWP at DWD. The horizontal grid nesting is addressed by two variants of the Jablonowski–Williamson test case (Jablonowski and Williamson, 2006): one considering the evolution of a baroclinic wave train triggered by a disturbance in the initial condition and one ideally remaining in steady state in the absence of grid irregularities. Choosing the global mesh size such that the baroclinic wave is poorly resolved, the former test shows that a refined grid passed by the wave during the early stage of its development has some beneficial impact on the quality of the simulation, deliver-

ing a result lying between non-nested reference runs with the respective coarse and fine mesh sizes. Numerical disturbances related to the resolution jump along the nest boundaries are clearly smaller than the benefit of the regionally refined resolution. The steady-state variant of the test nevertheless shows that this kind of disturbance does exist and is somewhat larger than the disturbances induced by the irregularities of the uniform global icosahedral grid. The accuracy of the vertical nesting is examined with the Schär mountain test case (Schär et al., 2002), which considers a

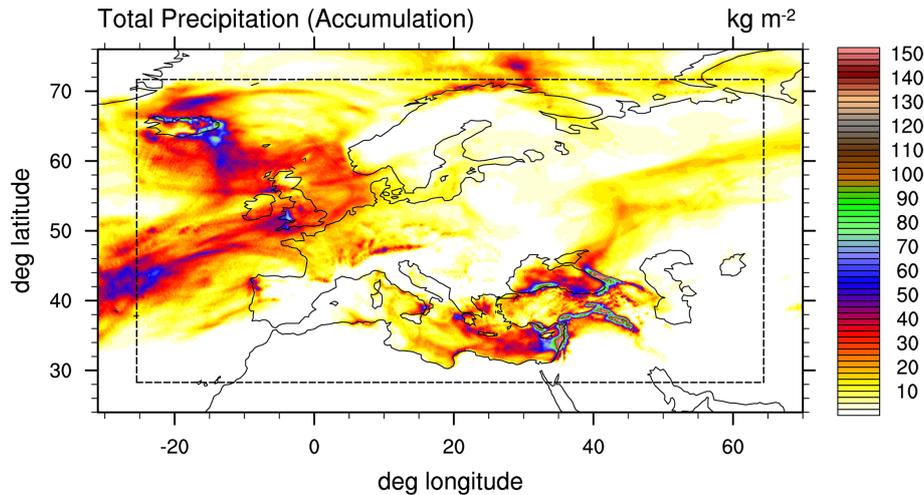


Figure 15. Accumulated total precipitation (kg m^{-2}) for a 5 d forecast starting at 00:00 UTC on 15 January 2021. The nest outline is indicated by the dashed box. Inside (outside) the dashed box, the solution of the nested (global) domain is shown.

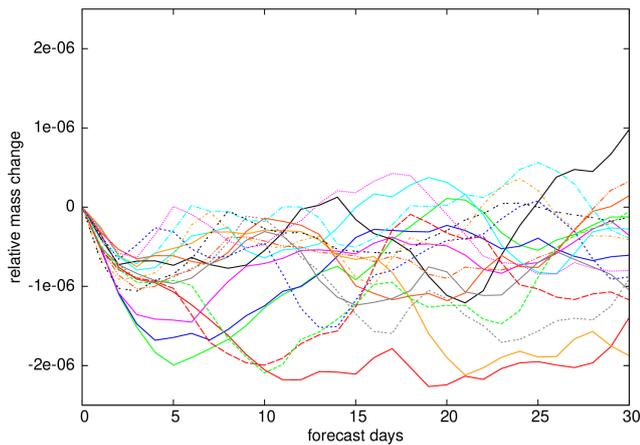


Figure 16. Relative mass conservation error in the global domain for 18 selected 30 d forecast experiments in January 2021 in which the nested domain remained active throughout the forecast.

steady-state multi-scale orographic gravity wave composed of vertically propagating and vertically decaying wave components. The results show that the wave reflections generated at the vertical nest interface are small enough to be irrelevant for practical applications. The findings from these idealized tests are corroborated by full-physics NWP experiments using the operational DWD domain configuration with a vertically nested grid over Europe and adjacent regions. Provided that the mesh sizes are chosen to lie in a range in which the forecast quality has a pronounced resolution dependence, these tests demonstrate that the feedback from the regionally refined domain to the global one has a significant beneficial impact on a variety of standard NWP scores, which also becomes evident downstream (over Asia) with a delay of about 3 d. We are thus able to conclude that the ICON grid nesting,

which has been used for operational weather forecasting at DWD since July 2015, has a practical and measurable benefit for our forecast quality.

Appendix A: Implementation aspects

A1 Grid point indexing

In order to identify grid points belonging to certain regions of a domain and to control the feedback between parent and child domains, dedicated integer fields named `refin_ctrl` exist for cells, edges and vertices (see Fig. 1). They are provided by ICON's grid generator. Along lateral nest boundaries, these fields contain positive numbers. For cells, they indicate the shortest distance to the boundary in units of cell rows. For example, a value of 1 indicates the outermost cells. A similar counting is applied for vertices. For edges, however, the counting proceeds twice as fast; i.e., the edge-based counter is the sum of two neighboring cell counters. This fine-grained counting enables the explicit specification of cell rows or edges up to which individual numerical operators (such as divergence, gradient or diffusion operators) are evaluated. By default, the 12 outermost cell rows (encompassing the boundary interpolation and feedback zones) are flagged by positive numbers. For the example in Fig. 1, ICON's grid generator had been configured to flag the 14 outermost cell rows, which then allows for an extended nudging zone of 10 cell rows at most.

In regions with an overlapping child domain the counters for cells, edges and vertices are filled with negative numbers (see light-blue area in Fig. 1). The counters start from -1 at the location of the child domain's outer boundary and descend towards the interior of the nest overlap region. The procedure is equivalent to the procedure along the lateral bound-

ary, with the exception that it already stops after three cell rows. All remaining interior points are flagged with a value of -4 for cells and vertices and -8 for edges. This flagging of cells overlapping a child domain allows easy access to all the points involved in the child-to-parent feedback process described in Sect. 2.2.2.

A2 Distributed-memory parallelization

Several measures are taken in order to optimize the computational efficiency of the nesting implementation.

Model grid points in ICON are stored in (blocked) 1D vectors. In order to achieve efficient runtime flow control without the need for masking operations within loops, the grid points are ordered by their distance from the lateral boundary by making use of the `refin_ctrl` fields described in Appendix A1. Hence, grid points lying at or near the lateral boundary of a nested domain are shifted to the beginning of the index vector. This allows excluding boundary points from prognostic computations accessing nonexistent neighbor points without masking operations. In the present implementation, the four outer cell rows constituting the boundary interpolation zone (Fig. 1) and the adjacent fifth one participate in the reordering. The remaining grid points follow in a non-ordered fashion, while the `refin_ctrl` indexing proceeds with the remaining grid points in the nudging zone, followed by the prognostic grid points not overlapping a child domain (`refin_ctrl = 0`), and by the overlapping prognostic points (`refin_ctrl < 0`) (see Fig. 1). In the presence of an MPI domain decomposition, the related halo points are shifted to the very end of the index vector. Moreover, the two outermost cell rows are counted with zero weight when computing the domain decomposition, implying that the benefit of these grid points doing very little computation is not lost by load imbalance.

Regarding distributed-memory (MPI) parallelization, the default strategy adopted in ICON is to distribute all model domains among all compute processors. As this implies that child grid points are in general owned by a different processor than the corresponding parent grid point, an intermediate layer having the mesh size of the parent grid but the domain decomposition of the child grid is inserted in order to accommodate the data exchange required for boundary interpolation and feedback.

To reduce the amount of MPI communication for complex nested configurations, multiple nested domains at the same nesting level can be merged into one logical domain, which is then not geometrically contiguous. This needs to be specified by the user during the grid generation process by indicating a list of individual domains that are supposed to be merged. The lateral boundary points belonging to all components of the merged domain are then collected at the beginning of the index vector. For all prognostic calculations, the multiple domains are treated as a single logical entity, and just the output files may be split according to the geometrically contiguous

basic domains. As one-way and two-way nesting cannot be mixed within one logical domain, there may still need to be two logical domains on a given nest level.

To further optimize the amount of MPI communication, a so-called processor splitting is available that allows for executing several nested domains concurrently on processor subsets whose size can be determined by the user in order to minimize the ensuing load imbalance. Unlike domain merging, this allows parallel execution of one-way and two-way nested domains. This option is currently restricted to the step from the global domain to the first nesting level in order to keep the technical complexity at a manageable level.

Code and data availability. The ICON release version icon-2.6.4 is freely available under a personal noncommercial research license. Information on the license and instructions for downloading the code can be found at https://code.mpimet.mpg.de/projects/iconpublic/wiki/Instructions_to_obtain_the_ICON_model_code_with_a_personal_non-commercial_research_license (MPI-M, 2019). By downloading the code the user accepts the license agreement. All primary data and scripts which are necessary to validate the research findings of Sect. 3.1 and 3.2 are freely available for download from the Open Research Data Repository of the Max Planck Society (Edmond) under <https://doi.org/10.17617/3.NOC2AE> (Zängl et al., 2022). Data and scripts for the NWP applications of Sect. 3.3 are not included due to the huge size of the data set (about 3.3 TB) and the fact that installing and running the verification toolchain outside DWD's computer systems is beyond the authors' expertise. Access to these data will be granted upon request (contact: guenther.zaengl@dwd.de).

Author contributions. All authors contributed to the writing of this paper. GZ conducted the initial nesting implementation and ran the experiments of Sect. 3.1 and Sect. 3.3. DR worked on the vertical nesting and ran the experiments of Sect. 3.2. FP developed the grid generator and worked on parallelization aspects and performance optimizations. Some elements of the description in Sect. 2 have already been published in the ICON user tutorial (Prill et al., 2020).

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Acknowledgements. The authors would like to thank the entire ICON team for their excellent work, without which this would not have been possible. Our additional thanks go to Felix Fundel (DWD), who developed the verification toolchain used in this work. The authors also appreciate the constructive comments made by the

two anonymous reviewers, which led to a significant improvement of the paper.

Review statement. This paper was edited by Ludovic Räss and reviewed by two anonymous referees.

References

- Bonaventura, L. and Ringler, T.: Analysis of discrete shallow water models on geodesic Delaunay grids with C-type staggering, *Mon. Weather Rev.*, 133, 2351–2373, <https://doi.org/10.1175/MWR2986.1>, 2005.
- Borchert, S., Zhou, G., Baldauf, M., Schmidt, H., Zängl, G., and Reinert, D.: The upper-atmosphere extension of the ICON general circulation model (version: ua-icon-1.0), *Geosci. Model Dev.*, 12, 3541–3569, <https://doi.org/10.5194/gmd-12-3541-2019>, 2019.
- Danilov, S.: On utility of triangular C-grid type discretization for numerical modeling of large-scale ocean flows, *Ocean Dynam.*, 60, 1361–1369, <https://doi.org/10.1007/s10236-010-0339-6>, 2012.
- Davies, H.: A lateral boundary formulation for multi-level prediction models, *Q. J. Roy. Meteor. Soc.*, 102, 405–418, <https://doi.org/10.1002/qj.49710243210>, 1976.
- Davies, T.: Lateral boundary conditions for limited area models, *Q. J. Roy. Meteor. Soc.*, 140, 185–196, <https://doi.org/10.1002/qj.2127>, 2014.
- Dubos, T. and Kevlahan, N. K.-R.: A conservative adaptive wavelet method for the shallow-water equations on staggered grids, *Q. J. Roy. Meteor. Soc.*, 139, 1997–2020, <https://doi.org/10.1002/qj.2097>, 2013.
- Fox-Rabinovitz, M., Cote, J., Dugas, B., Deque, M., McGregor, J. L., and Belochitski, A.: Stretched-grid Model Intercomparison Project: decadal regional climate simulations with enhanced variable and uniform-resolution GCMs, *Meteorol. Atmos. Phys.*, 100, 159–178, <https://doi.org/10.1007/s00703-008-0301-z>, 2008.
- Gassmann, A.: Inspection of hexagonal and triangular C-grid discretizations of the shallow water equations, *J. Comput. Phys.*, 230, 2706–2721, <https://doi.org/10.1016/j.jcp.2011.01.014>, 2011.
- Gassmann, A.: A global hexagonal C-grid non-hydrostatic dynamical core (ICON-IAP) designed for energetic consistency, *Q. J. Roy. Meteor. Soc.*, 139, 152–175, <https://doi.org/10.1002/qj.1960>, 2013.
- Gassmann, A. and Herzog, H.-J.: Towards a consistent numerical compressible non-hydrostatic model using generalized Hamiltonian tools, *Q. J. Roy. Meteor. Soc.*, 134, 1597–1613, <https://doi.org/10.1002/qj.297>, 2008.
- Gottelman, A., Callaghan, P., Larson, V. E., Zarzycki, C. M., Bacmeister, J. T., Lauritzen, P. H., Bogenschutz, P. A., and Neale, R. B.: Regional Climate Simulations with the Community Earth System Model, *J. Adv. Model Earth Sy.*, 10, 1245–1265, <https://doi.org/10.1002/2017MS001227>, 2018.
- Goto, D., Dai, T., Satoh, M., Tomita, H., Uchida, J., Misawa, S., Inoue, T., Tsuruta, H., Ueda, K., Ng, C. F. S., Takami, A., Sugimoto, N., Shimizu, A., Ohara, T., and Nakajima, T.: Application of a global nonhydrostatic model with a stretched-grid system to regional aerosol simulations around Japan, *Geosci. Model Dev.*, 8, 235–259, <https://doi.org/10.5194/gmd-8-235-2015>, 2015.
- Grell, G. A., Dudhia, J., and Stauffer, D.: A description of the fifth-generation Penn State/NCAR Mesoscale Model (MM5), Tech. Rep. NCAR/TN-398+STR, University Corporation for Atmospheric Research, <https://doi.org/10.5065/D60Z716B>, 1994.
- Gross, E. S., Bonaventura, L., and Rosatti, G.: Consistency with continuity in conservative advection schemes for free-surface models, *Int. J. Numer. Meth. Fl.*, 38, 307–327, <https://doi.org/10.1002/flid.222>, 2002.
- Gross, M., Wan, H., Rasch, P. J., Caldwell, P. M., Williamson, D. L., Klocke, D., Jablonowski, C., Thatcher, D. R., Wood, N., Cullen, M., Beare, B., Willett, M., Lemarié, F., Blayo, E., Malardel, S., Termonia, P., Gassmann, A., Lauritzen, P. H., Johansen, H., Zarzycki, C. M., Sakaguchi, K., and Leung, R.: Physics–Dynamics Coupling in Weather, Climate, and Earth System Models: Challenges and Recent Progress, *Mon. Weather Rev.*, 146, 3505–3544, <https://doi.org/10.1175/MWR-D-17-0345.1>, 2018.
- Harris, L. M. and Lin, S.-J.: A Two-Way Nested Global-Regional Dynamical Core on the Cubed-Sphere Grid, *Mon. Weather Rev.*, 141, 283–306, <https://doi.org/10.1175/MWR-D-11-00201.1>, 2013.
- Harris, L. M. and Lin, S.-J.: Global-to-Regional Nested Grid Climate Simulations in the GFDL High Resolution Atmospheric Model, *J. Climate*, 27, 4890–4910, <https://doi.org/10.1175/JCLI-D-13-00596.1>, 2014.
- Jablonowski, C. and Williamson, D.: A baroclinic instability test case for atmospheric model dynamical cores, *Q. J. Roy. Meteor. Soc.*, 132, 2943–2975, <https://doi.org/10.1256/qj.06.12>, 2006.
- Klemp, J. B., Skamarock, W. C., and Park, S.-H.: Idealized global nonhydrostatic atmospheric test cases on a reduced-radius sphere, *J. Adv. Model Earth Sy.*, 7, 1155–1177, <https://doi.org/10.1002/2015MS000435>, 2015.
- Lauritzen, P. H., Jablonowski, C., Taylor, M. A., and Nair, R. D.: Rotated Versions of the Jablonowski Steady-State and Baroclinic Wave Test Cases: A Dynamical Core Intercomparison, *J. Adv. Model Earth Sy.*, 2, 15, <https://doi.org/10.3894/JAMES.2010.2.15>, 2010.
- Leuenberger, D., Koller, M., Fuhrer, O., and Schär, C.: A generalization of the SLEVE vertical coordinate, *Mon. Weather Rev.*, 138, 3683–3689, <https://doi.org/10.1175/2010MWR3307.1>, 2010.
- Mouallem, J., Harris, L., and Benson, R.: Multiple same-level and telescoping nesting in GFDL's dynamical core, *Geosci. Model Dev.*, 15, 4355–4371, <https://doi.org/10.5194/gmd-15-4355-2022>, 2022.
- MPI-M: Instructions for obtaining the ICON Code, https://code.mpimet.mpg.de/projects/iconpublic/wiki/Instructions_to_obtain_the_ICON_model_code_with_a_personal_non-commercial_research_license, last access: 18 November 2019.
- Narcowich, F. and Ward, J.: Generalized Hermite interpolation via matrix-valued conditionally positive definite functions, *Math. Comput.*, 63, 661–687, <https://doi.org/10.2307/2153288>, 1994.
- Prill, F., Reinert, D., Rieger, D., and Zängl, G.: ICON Tutorial: Working with the ICON model, Deutscher Wetterdienst (DWD), https://doi.org/10.5676/dwd_pub/nwv/icon_tutorial2020, 2020.

- Reinert, D., Prill, F., Frank, H., Denhard, M., Baldauf, M., Schraff, C., Gebhardt, C., Marsigli, C., and Zängl, G.: DWD Database Reference for the Global and Regional ICON and ICON-EPS Forecasting System, Deutscher Wetterdienst (DWD), https://www.dwd.de/SharedDocs/downloads/DE/modelldokumentationen/nwv/icon/icon_dbbeschr_aktuell.html (last access: 15 August 2022), 2021.
- Rípodas, P., Gassmann, A., Förstner, J., Majewski, D., Giorgetta, M., Korn, P., Kornblueh, L., Wan, H., Zängl, G., Bonaventura, L., and Heinze, T.: Icosahedral Shallow Water Model (ICOSWM): results of shallow water test cases and sensitivity to model parameters, *Geosci. Model Dev.*, 2, 231–251, <https://doi.org/10.5194/gmd-2-231-2009>, 2009.
- Schär, C., Leuenberger, D., Fuhrer, O., Lüthi, D., and Girard, C.: A new terrain-following vertical coordinate formulation for atmospheric prediction models, *Mon. Weather Rev.*, 130, 2459–2480, [https://doi.org/10.1175/1520-0493\(2002\)130<2459:ANTFVC>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<2459:ANTFVC>2.0.CO;2), 2002.
- Schmidt, F.: Variable fine mesh in spectral global models, *Beitr. Phys. Atmos.*, 50, 211–217, 1977.
- Skamarock, W. C., Klemp, J. B., Duda, M. G., Fowler, L. D., Park, S.-H., and Ringler, T. D.: A Multiscale Nonhydrostatic Atmospheric Model Using Centroidal Voronoi Tessellations and C-Grid Staggering, *Mon. Weather Rev.*, 140, 3090–3105, <https://doi.org/10.1175/MWR-D-11-00215.1>, 2012.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Zhiquan, L., Berner, J., Wang, W., Powers, J. G., Duda, M. G., , Barker, D. M., and Huang, X.: A Description of the Advanced Research WRF Model Version 4, No. ncar/tn-556+str, National Center For Atmospheric Research, Boulder, CO, <https://doi.org/10.5065/1dfh-6p97>, 2019.
- Staniforth, A. N. and Mitchell, H. L.: A Variable-Resolution Finite-Element Technique for Regional Forecasting with the Primitive Equations, *Mon. Weather Rev.*, 106, 439–447, [https://doi.org/10.1175/1520-0493\(1978\)106<0439:AVRFET>2.0.CO;2](https://doi.org/10.1175/1520-0493(1978)106<0439:AVRFET>2.0.CO;2), 1978.
- Tomita, H.: A Stretched Icosahedral Grid by a New Grid Transformation, *J. Meteorol. Soc. Jpn. Ser. II*, 86A, 107–119, <https://doi.org/10.2151/jmsj.86A.107>, 2008.
- Wan, H., Giorgetta, M. A., Zängl, G., Restelli, M., Majewski, D., Bonaventura, L., Fröhlich, K., Reinert, D., Rípodas, P., Kornblueh, L., and Förstner, J.: The ICON-1.2 hydrostatic atmospheric dynamical core on triangular grids – Part 1: Formulation and performance of the baseline version, *Geosci. Model Dev.*, 6, 735–763, <https://doi.org/10.5194/gmd-6-735-2013>, 2013.
- Warner, T. T., Peterson, R. A., and Treadon, R. E.: A Tutorial on Lateral Boundary Conditions as a Basic and Potentially Serious Limitation to Regional Numerical Weather Prediction, *B. Am. Meteorol. Soc.*, 78, 2599–2618, [https://doi.org/10.1175/1520-0477\(1997\)078<2599:ATOLBC>2.0.CO;2](https://doi.org/10.1175/1520-0477(1997)078<2599:ATOLBC>2.0.CO;2), 1997.
- Weimer, M., Buchmüller, J., Hoffmann, L., Kirner, O., Luo, B., Ruhnke, R., Steiner, M., Tritscher, I., and Braesicke, P.: Mountain-wave-induced polar stratospheric clouds and their representation in the global chemistry model ICON-ART, *Atmos. Chem. Phys.*, 21, 9515–9543, <https://doi.org/10.5194/acp-21-9515-2021>, 2021.
- WMO: Manual on the Global Data-processing and Forecasting System: Annex IV to the WMO Technical Regulations, WMO-no. 485, World Meteorological Organization, Geneva, Switzerland, ISBN 978-92-63-10485-4, 2019.
- Zängl, G., Reinert, D., Rípodas, P., and Baldauf, M.: The ICON (ICOSahedral Nonhydrostatic) modelling framework of DWD and MPI-M: Description of the nonhydrostatic dynamical core, *Q. J. Roy. Meteor. Soc.*, 141, 563–579, <https://doi.org/10.1002/qj.2378>, 2015.
- Zängl, G., Reinert, D., and Prill, F.: Grid Refinement in ICON v2.6.4 (research data), Edmond [data set], <https://doi.org/10.17617/3.NOC2AE>, 2022.
- Zarzycki, C. M., Jablonowski, C., and Taylor, M. A.: Using Variable-Resolution Meshes to Model Tropical Cyclones in the Community Atmosphere Model, *Mon. Weather Rev.*, 142, 1221–1239, <https://doi.org/10.1175/MWR-D-13-00179.1>, 2014.