

Regional Superparameterization in a Global Circulation Model Using Large Eddy Simulations

Jansson, Fredrik; van den Oord, Gijs; Pelupessy, Inti; Grönqvist, Johanna H.; Siebesma, A. Pier; Crommelin, Daan

DOI

[10.1029/2018MS001600](https://doi.org/10.1029/2018MS001600)

Publication date

2019

Document Version

Final published version

Published in

Journal of Advances in Modeling Earth Systems

Citation (APA)

Jansson, F., van den Oord, G., Pelupessy, I., Grönqvist, J. H., Siebesma, A. P., & Crommelin, D. (2019). Regional Superparameterization in a Global Circulation Model Using Large Eddy Simulations. *Journal of Advances in Modeling Earth Systems*, 11(9), 2958-2979. <https://doi.org/10.1029/2018MS001600>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



RESEARCH ARTICLE

10.1029/2018MS001600

Regional Superparameterization in a Global Circulation Model Using Large Eddy Simulations

Key Points:

- We present an efficient implementation of a large eddy simulation (LES)-based superparameterization in a global circulation model
- Superparameterization can be enabled for grid points in a selected area
- The initial results of an LES-based superparameterized simulation are evaluated and found promising

Supporting Information:

- Supporting Information S1-S3
- Movie S4
- Movie S5
- Movie S6

Correspondence to:

F. Jansson,
jansson@cwil.nl

Citation:

Jansson, F., van den Oord, G., Pelupessy, I., Grönqvist, J. H., Siebesma, A. P., & Crommelin, D. (2019). Regional superparameterization in a global circulation model using large eddy simulations. *Journal of Advances in Modeling Earth Systems*, 11, 2958–2979. <https://doi.org/10.1029/2018MS001600>

Received 21 DEC 2018

Accepted 6 JUL 2019

Accepted article online 12 JUL 2019

Published online 12 SEP 2019

Fredrik Jansson¹, Gijs van den Oord², Inti Pelupessy², Johanna H. Grönqvist³, A. Pier Siebesma^{4,5}, and Daan Crommelin^{1,6}

¹Centrum Wiskunde & Informatica, Amsterdam, Netherlands, ²Netherlands eScience Center, Amsterdam, Netherlands, ³Physics, Faculty of Science and Engineering, Åbo Akademi University, Turku, Finland, ⁴Royal Netherlands Meteorological Institute, ⁵Faculty of Civil Engineering and Geosciences, Delft University of Technology, Delft, Netherlands, ⁶Korteweg-de Vries Institute, University of Amsterdam, Amsterdam, Netherlands

Abstract As a computationally attractive alternative for global large eddy simulations (LESs), we investigate the possibility of using comprehensive three-dimensional LESs as a superparameterization that can replace all traditional parameterizations of atmospheric processes that are currently used in global models. We present the technical design for a replacement of the parameterization for clouds, convection, and turbulence of the global atmospheric model of the European Centre for Medium-Range Weather Forecasts by the Dutch Atmospheric Large Eddy Simulation model. The model coupling consists of bidirectional data exchange between the global model and the high-resolution LES models embedded within the columns of the global model. Our setup allows for selective superparameterization, that is, for applying superparameterization in local regions selected by the user, while keeping the standard parameterization of the global model intact outside this region. Computationally, this setup can result in major geographic load imbalance, because of the large difference in computational load between superparameterized and nonsuperparameterized model columns. To resolve this issue, we use a modular design where the local and global models are kept as distinct model codes and organize the model coupling such that all the local models run in parallel, separate from the global model. First simulation results, employing this design, demonstrate the potential of our approach.

1. Introduction

An accurate representation of clouds and convection in global weather and climate models and their interaction with the large-scale circulation remains one of the main challenges in atmospheric modeling. Uncertainties in the representation of clouds and convection are the prime sources of uncertainty in climate model sensitivity and major contributors to longstanding biases in the representation of the precipitation patterns in current climate and their projections in future climate (Bony & Dufresne, 2005; Bony et al., 2015; Schneider et al., 2017).

Cloud-related processes occur over a wide range of scales ranging from cloud droplet formation at the micrometer scale to cloud convective updrafts and downdrafts that can be as large as 10 km, from which organized mesoscale systems can emerge extending over hundreds of kilometers. Current operational global models operate at numerical resolutions in the range of 10–100 km. As a consequence, cloud and convective processes are numerically not resolved, and their impact on the resolved state is instead approximated by parameterizations, causing uncertainties of these unresolved processes.

The problem of parameterized clouds and convection is largely avoided when using large eddy simulations (LESs). The paradigm of LES is based on the idea that small unresolved turbulent eddies can be faithfully parameterized in terms of the resolved large eddies by making use of the self-similar structure of turbulence in the inertial subrange. The atmospheric inertial subrange is bounded by the depth of the atmospheric boundary layer, which has a typical depth of 1 km, indicating that a minimum resolution of $O(100\text{ m})$ is required to numerically resolve the relevant turbulence, convection, and cloud dynamics.

Two obvious but challenging pathways for improving the representation of clouds and convection in global models are either increasing the resolution of existing global models to turbulence-resolving scales or extending the spatial domains of LESs until a global scale is reached. Regarding the latter approach, realistic multiday LESs have been reported on domains approaching $1,000\text{ km}^2$ (Heinze et al., 2017; Schalkwijk

©2019. The Authors.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

et al., 2015). Global multiday simulations are approaching the 1-km resolution scale and are capable of partly resolving the cloud dynamics and therefore usually referred to as global cloud resolving models (CRMs) (Bretherton & Khairoutdinov, 2015; Miyamoto et al., 2013). Global CRMs form an interesting playground to explore the interaction between the global circulation and the resolved moist convective systems, but one should also bear in mind that for resolutions used in global CRMs (1–5 km), atmospheric turbulence and boundary layer clouds remain essentially unresolved.

A different pathway is offered by superparameterization (SP; Khairoutdinov & Randall, 2001; Khairoutdinov et al., 2005; Grabowski & Smolarkiewicz, 1999; Grabowski, 2001, 2004), where existing global and local CRMs are coupled. The common setup for SP is to replace deep convection and cloud parameterization schemes in every model column of the global model by a CRM. Because of computational constraints, the CRMs used in SP are mostly two-dimensional (2-D). A 3-D CRM was used by Khairoutdinov et al. (2005), but the grid of the CRM was still very coarse and limited to 8×8 columns. Jung and Arakawa (2010) present a quasi-3-D SP where the global models grid points are connected by narrow corridors consisting of 3-D local models. An overview of other simulation studies with SP is provided by Tao and Chern (2017). Like global CRMs, however, these SP approaches traditionally use horizontal resolution of 1–4 km and coarse vertical resolutions and still require additional parameterizations for boundary layer clouds and turbulence.

Recently, a variation on SP has been proposed (Grabowski, 2016), where the resolution of local CRMs becomes fine enough to be turbulence resolving. Parishani et al. (2017) included LES models with a fine spatial horizontal resolution (250 m) and 125 vertical levels in their global model. However, to be able to run the global model with SP, the LESs used small domains (8×8 columns, i.e., $2 \text{ km} \times 2 \text{ km}$) covering only a small fraction of the domain of a single column of the global model.

Ideally, SP is carried out with a 3-D high-resolution CRM that covers the full domain of each global model column. To reduce the enormous computational cost of this (hypothetical) SP setup, in the studies mentioned above, either the 3-D CRM is simplified to 2-D or quasi-3-D or the grid of each CRM is kept small (e.g., 8×8 horizontal). In the latter case, one can choose between high resolution on a small CRM domain (Parishani et al., 2017) and coarser resolution on a larger CRM domain (Khairoutdinov et al., 2005).

The SP approach that we will present in this study is different. Our aim is to use turbulence resolving resolutions on sufficiently large 3-D domains as a SP, in accordance with the resolution of the large-scale model. Computationally, this approach obviously does not allow SP to be applied globally yet. Therefore, rather than reducing the cost of the 3D CRM as sketched above, our setup creates the possibility to use SP only in a selected region, while leaving the regular (non-SP) parameterization in use outside this region. This way it is possible to explore the added value of superparameterized LES, albeit on a limited domain.

The benefits of such a 3-D LES-based SP in a conventional hydrostatic global model over a global LES or CRM have been discussed in Grabowski (2016). Computationally, it is attractive since all the local models can run independent from each other and only have to exchange mean profiles with the large-scale model. This allows an efficient implementation on massively parallel computer systems since all the SP models can run independently on separate cores. Further acceleration can be achieved by running the SP models sparser in space and time (Jones et al., 2015; Xing et al., 2009) or by varying the vertical resolution of the SP models depending on cloud types that they need to resolve (Marchand & Ackerman, 2011). Conceptually, there is also the advantage that the large-scale model can be formulated efficiently in a hydrostatic manner while the smaller scale LES-based SP models can be conveniently expressed in an anelastic formulation. This way there is no need to find an appropriate soundproof compressible formulation of the dynamics on a global scale. A major advantage compared to conventional CRM-based SP formulations (where the CRMs are coarse and therefore still contain parameterizations) is that with 3-D LES-based SP no parameterization for boundary layer clouds is required anymore.

The drawback of any SP formulation is that it introduces a scale break at the resolution scale of the large-scale model. This hinders the spectral transfer of variability across this scale and will potentially influence mesoscale organization. However, it also offers an excellent opportunity to explore the effect of such a scale break, which is present in every operational large-scale model.

In this paper we will discuss the implementation and performance of the Dutch Atmospheric Large Eddy Simulation (DALES; Heus et al., 2010) model as a regional 3-D LES-based SP into the Open Integrated Forecast System (OpenIFS) developed at the European Centre for Medium-Range Weather Forecasts (ECMWF)




Figure 1. Overview of the superparameterized model. Some grid columns of the global model OpenIFS (purple) are selected for superparameterization. Each of them is coupled to a local DALES model (blue), which resolves clouds and convection in three dimensions. The tendencies generated by these processes are fed back to the global model.

(Carver & Vana, 2017; ECMWF, 2014a, 2014b). We note that the focus of this paper is on presenting the design and implementation of our SP setup. We show some numerical results to demonstrate the feasibility of the approach but leave a more extensive validation to another study.

Section 2 describes the complete methodology of the coupling while section 3 concentrates more on the technical implementation. Section 4 presents results of a superparameterized atmospheric simulation over the Netherlands, comparisons with observations and an evaluation of the numerical performance. Section 5 contains conclusions, discussion, and outlook.

2. Methods

For coupling the global and CRMs, we follow the approach presented by Grabowski (2004) and also Benedict and Randall (2009), Khairoutdinov and Randall (2001), and Khairoutdinov et al. (2005). In the grid columns of the global model that are selected for SP, a local LES model is embedded as shown in Figure 1. The general idea is that for each coupled quantity, a forcing is introduced, which keeps the states of the two models consistent with each other. The coupling is bidirectional, so that the effects of clouds, turbulence, and convection, which are resolved in the local model, are fed back to the global model.

Below, we summarize the coupling procedure. In section 2.1 we discuss a simplified case where the two models are assumed to have similar vertical grid levels and to be formulated in terms of the same quantities, closely following the setup presented by Grabowski (2004). In the remaining parts of section 2 we discuss the adaptations needed to couple models with different vertical levels and different physical quantities, the surface model, and other modeling choices made in our setup with coupling between DALES and OpenIFS.

2.1. Physical Coupling of the Models

We consider a 3-D small-scale model embedded in a single column of the large-scale model. In earlier versions of the SP scheme (Grabowski & Smolarkiewicz, 1999; Grabowski, 2001), the states of the large-scale and the small-scale models are relaxed toward each other, with a freely chosen time constant, which was taken as 1 hr. In the later scheme (Grabowski, 2004), the relaxation time constant is chosen equal to




Figure 2. In OpenIFS the column physics schemes for turbulence, convection, and clouds are turned off in the superparameterized columns. These processes are instead handled by DALES.

the time step of the large-scale model, making the models more tightly coupled. The aim is that for any coupled quantity Q , the horizontal slab average in the small-scale model at height z matches the value at the same height in the large-scale model:

$$Q(Z = z, t) = \langle q(x, y, z, t) \rangle. \quad (1)$$

Capital letters denote quantities in the large-scale model; lower-case letters denote the small-scale model. Q and q here may represent any of the prognostic variables, and the brackets $\langle \cdot \rangle$ denote a horizontal slab average over the domain of the local model. In the appendix we analyze to what extent the desired equality (1) can be achieved with the scheme from Grabowski (2004).

The coupled variables generally include the horizontal wind velocities, the temperature, and the specific humidity. As in earlier SP works, the vertical velocities are left uncoupled. Since each local model is a closed system due to periodic boundary conditions, the horizontal average of the vertical velocity is zero.

The models are coupled by introducing additional forcings in both models, that is, additional contributions to the time derivatives of the coupled quantities. F_Q represents a forcing that stems from q and acts on Q in the large-scale model, while f_q represents a forcing stemming from Q and acting on q in the small-scale model.

The time stepping procedure is as follows. The large-scale model performs a single time step from time T to $T + \Delta T$, then the small-scale model is evolved over the same time interval, in multiple steps of length Δt . Before the time evolution of each model, forcings are calculated based on the difference between the most recently obtained states of the two models.

- (i) Given the state of both models at time T , represented by $Q(T)$ and $q(T)$, calculate forcings on the large-scale model

$$F_Q(T) = \frac{\langle q(T) \rangle - Q(T)}{\Delta T}. \quad (2)$$

- (ii) Time step the large-scale model

$$Q(T + \Delta T) = Q(T) + \Delta T[A_Q(T) + S_Q(T) + F_Q(T)], \quad (3)$$

where $A_Q(T)$ represents advection terms and $S_Q(T)$ represents source terms during the step from T to $T + \Delta T$.

- (iii) Calculate the forcing on the small-scale model

$$f_q(T) = \frac{Q(T + \Delta T) - \langle q(T) \rangle}{\Delta T}. \quad (4)$$

- (iv) Time step the small-scale model

$$q(T + \Delta T) = q(T) + \sum_{t=T}^{T+\Delta T} \Delta t [a_q(t) + s_q(t) + f_q(T)]. \quad (5)$$

The sums over t here schematically represent evolving the small-scale model over several time steps, with $a_q(t)$ denoting advection terms and $s_q(t)$ denoting source terms in the small-scale model.

This choice of forcings is such that they couple the advection and source terms between the models; see also Grabowski (2004). In particular, one can show that

$$F_Q(T + \Delta T) = \frac{1}{\Delta T} \sum_{t=T}^{T+\Delta T} \Delta t \langle a_q(t) + s_q(t) \rangle \quad (6)$$

and

$$f_q(T) = A_Q(T) + S_Q(T), \quad (7)$$

so that the forcings on the small-scale model equals the advective and source tendencies in the large-scale model and vice versa. Thus, each physical process should be taken into account in one of the models, but not in both. Otherwise, the contribution will be counted twice. As is shown in the appendix, the equality (1) is satisfied exactly if all physical processes are accounted for in one model and none in the other: If $\langle a_q(t) + s_q(t) \rangle = 0$ for all t , then $Q(T) = \langle q(T) \rangle$, whereas if $A_Q(T) + S_Q(T) = 0$, then $Q(T + \Delta T) = \langle q(t) \rangle$.

The Grabowski scheme does not explicitly describe the SP procedure for the sequential-splitting method in the global model, which is used in the column physics routines in OpenIFS. In this algorithm the physics processes are ordered by decreasing time scales, and every tendency is calculated with updated fields as its input, so that the tendencies of slower processes contribute to the evaluation of tendencies due to faster processes. We preserve this procedure by inserting the coupling to the local models at the stage in the OpenIFS time step where the parameterizations we substitute are evaluated, namely, turbulence, convection, and cloud schemes as shown in Figure 2.

2.2. Interpolation and Change of Variables

The coupling scheme outlined so far is the standard SP scheme as described in the references—where the vertical grids are assumed to be the same in the two models and the models are formulated in the same variables. In our case of coupling OpenIFS and DALES, neither of these assumptions can be made, requiring a few extra steps in the model coupling. The two models may use different vertical grids; typically, the local model has a denser grid than the global model and will not extend beyond the tropopause. OpenIFS is formulated in so-called hybrid sigma pressure coordinates; in our test cases we used 90 vertical levels extending up to roughly 80 km. For DALES we typically used a vertical spacing of 25 m, extending up to 4 km. We note that DALES was initially formulated with the Boussinesq approximation (Heus et al., 2010). The anelastic approximation, which makes the model suitable for larger vertical extents, was introduced in DALES 4.0 in 2014 (see, e.g., Böing et al., 2012).

To exchange vertical profiles of quantities between global and the local models, we use linear interpolation along the z axis. We convert the OpenIFS hybrid model level profiles to altitude by fetching the full- and half-level pressure profiles P_f, P_h and using the hydrostatic approximation to determine the height of each layer:

$$dz = \frac{-R_d T [1 + (R_v/R_d - 1)Q_v - (Q_L + Q_I)]}{g_0 P_f} dP. \quad (8)$$

Here, $R_d \approx 287.04$ J/kg K is the gas constant for dry air, $R_v \approx 461.5$ J/kg K is the gas constant for water vapor, and g_0 is the acceleration due to gravity. Q_v is the water vapor specific humidity, while Q_L and Q_I are the specific humidities of liquid water and ice. Above the vertical extent of the DALES models, we set the forcings F_Q on the global model to zero.

Furthermore, the two models are formulated using different prognostic variables. The coupling thus requires a variable conversion step. Table 1 lists the quantities that are coupled between the two models; the conversions required are described in detail below.

For temperature, OpenIFS uses a regular temperature T , while DALES uses the liquid water potential temperature θ_l (Heus et al., 2010). For the conversion, we use

$$\theta_l \approx \frac{T}{\Pi} - \frac{L}{c_{pd}\Pi} q_c, \quad (9)$$

where q_c is the cloud condensate defined as the sum of cloud liquid water q_L and cloud ice q_I . The Exner function Π is defined as

$$\Pi(p) = \left(\frac{p}{p_0} \right)^{R_d/c_{pd}}. \quad (10)$$

Here $c_{pd} \approx 1,004$ J/kg K is the specific heat of dry air at constant pressure.

Table 1
Quantities in OpenIFS and DALES used in the Superparameterization Scheme

OpenIFS quantities				
Symbol	Internal array name	Description	Unit	Coupling direction
U, V	GMV	Horizontal velocity	m/s	Bidirectional
T	GMV	Temperature	K	Bidirectional
Q_V	GFL	Specific water vapor content	kg/kg	Input ^a
Q_L	GFL	Specific liquid water content	kg/kg	Input ^a
Q_I	GFL	Specific cloud ice content	kg/kg	Input ^a
Q_T	$Q_V + Q_L + Q_I$	Total specific humidity	kg/kg	Output ^a
Z_{0M}, Z_{0H}	ZAZ0M, ZAZ0H	Surface roughness	m	Output
F_{SH}	PDIFTQ	Specific humidity flux	kg/m ² s	Output
F_{QL}	PDIFTL	Specific cloud liquid water flux	kg/m ² s	Output
F_{QI}	PDIFTI	Specific cloud ice flux	kg/m ² s	Output
F_{TS}	PDIFTS	Sensible heat flux	W/m ²	Output
P_S	GMVS	Surface pressure	Pa	Output
A	GFL	Cloud fraction	—	Input
DALES quantities				
Symbol	Variable name	Description	Unit	Coupling direction
u, v	u0, v0	Horizontal velocity	m/s	Bidirectional
θ_l	thl0	Liquid water potential temperature	K	Bidirectional
q_t	qt0	Specific humidity	kg/kg	Input ^a
q_v		Specific water vapor content	kg/kg	Output ^a
q_L		Specific cloud liquid water content	kg/kg	Output ^a
q_I		Specific cloud ice content	kg/kg	Output ^a
F_q	wqsurf	Surface kinematic moisture flux	kg/kg m/s	Input
F_{θ_l}	wtsurf	Surface kinematic θ_l flux	K m/s	Input
P_S		Surface pressure	Pa	Input
A		Cloud fraction	—	Output

^aHumidity requires a special treatment due to the different choices of prognostic variables in OpenIFS and DALES: The total humidity q_t in DALES is forced toward Q_T , which is calculated in OpenIFS as the sum of the vapor, liquid, and ice humidities. The OpenIFS quantities Q_V , Q_L , and Q_I are forced toward the corresponding diagnosed quantities q_v , q_L , and q_I in DALES. See also section 2.2.

For the humidity, DALES uses only a total (nonprecipitating) humidity q_t as a prognostic variable. At every time step, q_t is partitioned into vapor, cloud liquid water, and cloud ice according to the local temperature and pressure. (Humidity above the saturation humidity q_{sat} is assumed to convert instantly into cloud condensate, q_c . The condensate is split between liquid and ice based on temperature and goes linearly from all liquid when $T \geq 268$ K to all ice when $T \leq 253$ K.)

OpenIFS, on the other hand, has separate prognostic variables for these quantities, Q_V for water vapor, Q_L for liquid water, and Q_I for ice. When calculating forcings on DALES, the total humidity q_t is nudged toward the total humidity of the global model, calculated as $Q_T = Q_V + Q_L + Q_I$. When coupling the humidities back to the global model, the diagnosed values of the DALES are used, so that each one of Q_V , Q_L , and Q_I is forced toward the horizontal average of the corresponding diagnosed quantity in DALES.

2.3. Surface Scheme

Both DALES and OpenIFS contain a surface model, which accounts for surface drag and for fluxes of heat and moisture between the atmosphere and the land or sea surface. We have chosen to use fluxes and surface roughness lengths calculated in OpenIFS, while letting DALES handle the effects these have on the atmosphere at the superparameterized grid points. DALES does not explicitly represent the surface topography. The topography on scales smaller than the global model grid is represented by the surface roughness length.




Figure 3. Organization of the superparameterized simulation code, with a top-level coupling program communicating with the global model and a number of local model instances through OMUSE interfaces.

In this way, we can rely on the land/sea mask, soil and vegetation data, and ocean wave model of OpenIFS, making it easy to set up a SP anywhere without having to supply detailed surface information.

We achieve this by having DALES run with prescribed roughness lengths and surface fluxes of moisture and heat. These quantities are retrieved from the OpenIFS at every time step. DALES includes a surface scheme based on Monin-Obukhov similarity theory, which calculates surface drag given the surface roughness length. To avoid double counting of the surface fluxes, we disable the contribution of the surface layer scheme in OpenIFS at the superparameterized grid points. This procedure provides a simple way to account for orographic variations and vegetation for the local models located over land, and for wave height for models above the sea.

Since OpenIFS and DALES are built using quantities with different units (see Table 1), unit conversions are necessary to consistently couple the surface fluxes. For the humidity flux, a scaling with the air density ρ is required:

$$\mathcal{F}_q = -\frac{\mathcal{F}_{QL} + \mathcal{F}_{QI} + \mathcal{F}_{SH}}{\rho}. \quad (11)$$

In the case of the surface heat flux, the conversion from \mathcal{F}_{TS} in OpenIFS to \mathcal{F}_{θ_1} in DALES is

$$\mathcal{F}_{\theta_1} = -\frac{\mathcal{F}_{TS}}{\Pi(P_S)c_{pd}\rho}, \quad (12)$$

where Π is the Exner function given in equation (10). We have used the air density ρ from OpenIFS in these expressions (the densities in the two models are generally very close due to the model coupling). The models also differ in sign conventions: In DALES positive fluxes are upward, into the atmosphere, while in OpenIFS positive fluxes are downward.

2.4. Radiation, Cloud Condensate, and Cloud Fraction

In our setup, radiative heating and cooling in the atmosphere are currently handled only in the global model. The radiation scheme in OpenIFS uses the cloud fraction A and the cloud condensate content Q_L , Q_I as input quantities. Since these quantities are coupled to the local models in the superparameterized columns, the radiation scheme operates on a state that is consistent with the local model but at a coarser resolution. Furthermore, the OpenIFS scheme is only active once every few time steps to save compute time.

The cloud fraction A in the global model is derived by calculating the fraction of all the columns in DALES that contain a nonzero cloud condensate in the range from k_1 to k_2 . In formula, the cloud fraction A in the global model is given by

$$A = \frac{1}{i_{\max}j_{\max}} \sum_{i=1}^{i_{\max}} \sum_{j=1}^{j_{\max}} I^{k_2, k_1}(i, j), \quad (13)$$

where i and j are the horizontal grid indices of DALES and I is an indicator function that takes the value 1 in the case of any cloud condensate in the subcolumn between levels k_1 and k_2 at the horizontal coordinates i and j , and zero otherwise (Neggers et al., 2011). This choice for deriving A as a cloud fraction “defined by area,” or projected cloud fraction is deliberate, because in OpenIFS it is implicitly assumed that clouds do not exhibit any subgrid variability within the vertical extent of the layer. Physically, this implies that in the case where not all vertical levels in a DALES subcolumn are occupied with cloud condensate, it is averaged out over this subcolumn.

In principle, one could use the DALES radiation scheme to obtain more realistic heating and cooling rates in cloudy layers: DALES supports the rapid radiative transfer model that is similar to the OpenIFS scheme. Our choice to switch it off is motivated by performance considerations, as the DALES radiation scheme will

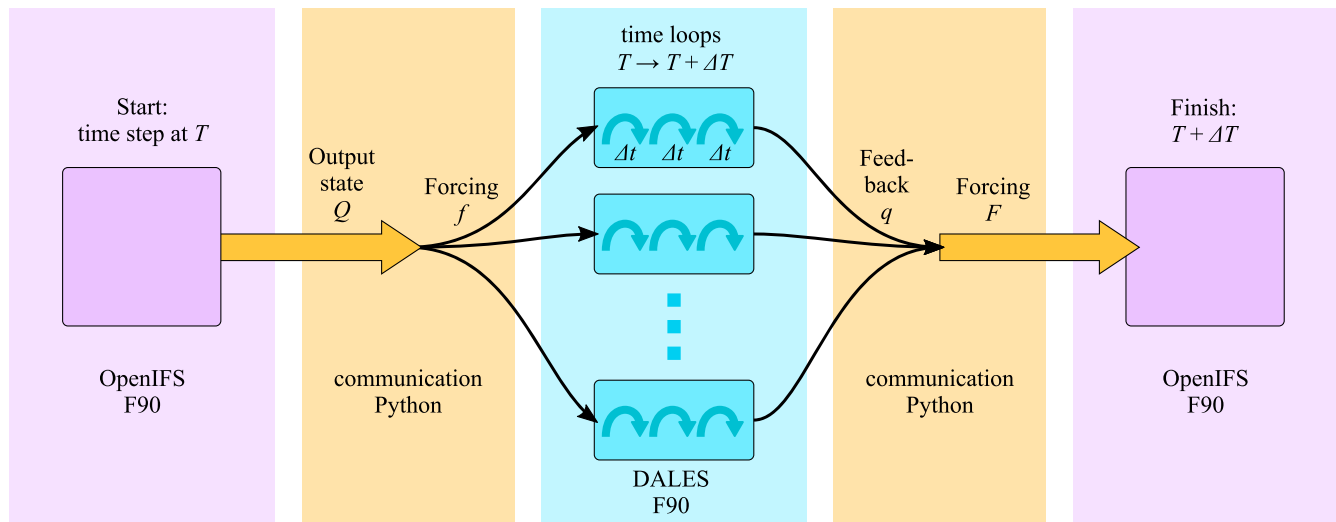


Figure 4. Single time step of the superparameterized OpenIFS.

make the local time stepping significantly slower, and technical considerations, concerning the atmosphere above the DALES instances. In Rapid Radiative Transfer Model, a system of equations is solved to obtain the radiation flux upwards and downwards at each model layer, accounting for scattering and absorption. As the set of equations is solved at once for the full vertical range of the atmosphere, it is difficult to separate the radiation scheme into separate parts handled in the two models.

The straight-forward but expensive solution would be to extend the local model to cover the same vertical range as the global model, so that it can handle the complete radiation calculations. In previous SP work, it has been a common choice to use the same vertical levels in both the local and the global models.

An alternative scheme where the LES model can have a smaller vertical extent than the global model is the following: Both DALES and OpenIFS perform their own radiation calculations. DALES performs radiation calculations at its own high spatial and temporal resolution. For the part of the atmosphere above the simulation volume, humidity, and trace gas profiles from OpenIFS are used. The resulting radiative heating tendencies are applied in the DALES model. In OpenIFS, radiation is handled as before, at the coarse spatial and temporal resolution. However, the resulting tendencies are set to 0 for the model layers overlapping the DALES model, to avoid counting the radiation tendencies twice (see section 2.1).

As a final remark, for full benefits of the high-resolution radiation scheme, also, the surface should be modeled at the higher resolution of the LES, so that the surface fluxes would be consistent with the clouds present. We expect this to be more important over land, where the surface temperature reacts rapidly to changes in the incident radiation.

2.5. Other Modeling Choices

In DALES, the highest 25% of the model levels constitutes a so-called sponge layer. It removes fluctuations of wind, temperature, and humidity, in order to damp gravity waves before they can reflect at the model top. The damping increases smoothly from the start of the sponge layer to the top of the system. There are several options for the damping mechanism. To be compatible with the SP, we use a scheme where each quantity is relaxed toward the horizontal average of that quantity. Since this scheme preserves the horizontal average, it does not have a strong effect on the global model.

As in earlier SP schemes, we do not couple the vertical velocity w between the two models. Due to the periodic boundary conditions of the local model, the horizontal average of $\langle w \rangle$ vanishes, which excludes the possibility of an explicit coupling. The effect of the large-scale vertical velocity on the prognostic fields in the global model is of course taken into account by the vertical advection, and its effect is felt by the local model through the forcing as expressed by equation (4).

Table 2
Parameters Used in the Simulations

Coupling parameters	
Number of DALES instances	42
Spin-up time	4 hr
Duration	21 hr
DALES parameters	
Vertical resolution	25 m
Horizontal resolution	200 m
Vertical extent	4 km
Horizontal extent	40 km
Grid size	200 × 200 × 160
Governing equations	Anelastic approximation
Turbulence scheme	subfilter-scale turbulence kinetic energy SFS-TKE ^a
Advection scheme - vertical	2nd order ^a
Advection scheme - horizontal	6th order (momentum), 5th order nearly monotonous (scalars) ^a
Time stepping scheme	3rd-order Runge-Kutta ^a
Time step	Adaptive, ≈ 1 ... 20 s ^a
Time step limit	$\Delta t \sqrt{(u/\Delta x)^2 + (v/\Delta y)^2 + (w/\Delta z)^2} < 0.7$
Time step diffusion limit	$\Delta t D (1/\Delta x^2 + 1/\Delta y^2 + 1/\Delta z^2) < 0.3$
OpenIFS parameters	
Grid	T511L91
Grid point distance	≈40 km
Vertical resolution at ground	≈23 m
Vertical resolution at 1 km	≈163 m
Vertical resolution at 4 km	≈356 m
Initial state and sea surface temperature forcing	ERA-Interim
Start time	2012-04-13 at 00:00 UTC
Time step	900 s

^aFor details of the DALES model, see Heus et al. (2010), as well as Wicker and Skamarock (2002) for details of the numerical schemes.

2.6. DALES Horizontal Extent and Resolution

Since SP does not involve lateral boundary forcings and all exchanged profiles involve bulk properties, the horizontal extent of the local models can in principle be chosen independently from the OpenIFS grid spacing. The DALES model size should be chosen to capture mesoscale structures as much as possible. However, the occurrence of such organization very much depends on the interaction with the large-scale dynamics and is often difficult to predict. For the horizontal DALES grid spacing we use 200 m; beyond this scale, the DALES subgrid model can no longer accurately account for the unresolved turbulent motions.

3. Implementation

For SP the OpenIFS model needs a bidirectional coupling to multiple instances of DALES, each one mapping to a different grid point of the global model. The central hypothesis in the software design is that the bulk of the computing time in the coupled system is going to be spent during the time stepping of the local models. This is primarily because OpenIFS has a coarse grid and its numerical scheme allows large time steps, whereas the three-dimensional small-scale models are frequently restricted in their time step and thus have to perform many iterations to catch up with the global model. Hence, it is paramount to allow the independent DALES instances to run concurrently on separate resources whenever available.

Previous SP setups have generally embedded the local model in the column physics routines of the global model. This approach allows for computationally efficient parallelization when (i) SP is applied uniformly

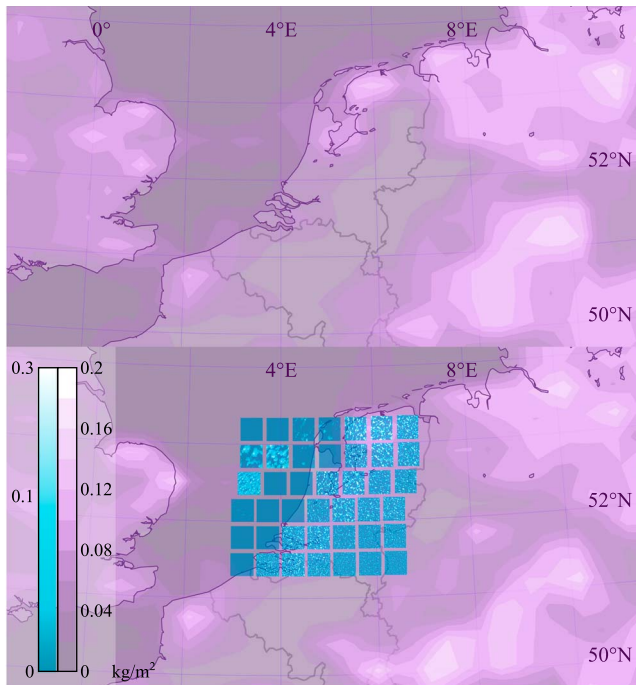


Figure 5. A superparameterized simulation over the Netherlands (bottom), compared to the same simulation run in standard OpenIFS (top). DALES instances are shown in blue, over a background showing the OpenIFS state in purple. Cloudiness (liquid water path) is shown in shades of white for both models. The state shown is for the time 13 April 2012, 11:35 UTC; the simulation was started at midnight the same date, with a state from ERA-Interim.

in all model columns, (ii) the global model has a balanced grid partitioning, and (iii) the local model compute time does not vary much between grid boxes. In that case all parallel processes (e.g., Message Passing Interface (MPI) tasks) carry a similar burden of computing the superparameterizing tendencies.

In a regional SP, however, this approach would lead to an extreme imbalance in the compute load between parallel subdomains containing superparameterized columns and those relying only on the traditional fast parameterizations. Therefore, it is necessary to redistribute the DALES computational work over the available compute resources. In fact, it suffices to have the majority of the compute resources dedicated to pure DALES time stepping alone and let the global model run on a small subset of cores. Since managing the allocation of resources per component is essential for achieving an efficient coupled system, we have chosen to design our model coupling as a master coupling code, which addresses the global and local models as separate library components.

The coupling program is written in Python and communicates with the models, which are written mainly in Fortran, through a framework named OMUSE. (Pelupessy et al., 2017, see also ; Pelupessy et al., 2013; Portegies Zwart et al., 2009; Portegies Zwart et al., 2013). This setup is illustrated in Figure 3. In the following sections, we explain the modifications made in OpenIFS and DALES, the role of the OMUSE framework, and the coupling program.

3.1. Interface to OpenIFS

For the SP coupler to be able to communicate with OpenIFS, an interface has been defined, with functions for initializing the model, setting tendencies, performing a time step, and retrieving the model state.

In the original OpenIFS physics routines, the different physical processes are evaluated in sequence, in order of decreasing characteristic time scale. The model uses a so-called sequential-splitting time stepping scheme, where the model state is updated after each process, so that later processes operate on a state modified by the earlier ones. To preserve this process ordering in the superparameterized model, the coupling to the small-scale model should take place at the stage of replaced processes, namely, where the (boundary layer) turbulence, the convection, and the cloud scheme in the OpenIFS physics routine are called, as shown in Figure 2.

The vertical physics processes in the original OpenIFS are evaluated within each MPI task of OpenIFS, for one column at a time, in a sequential horizontal loop (actually, the physical processes act on blocks of columns for optimal cache usage and vectorization). With this architecture, the loop would be halted at every SP column, awaiting the completion of the local model, before moving to the next column in the loop. To overcome this problem, we restructured the time loop in OpenIFS and split it into three parts: (i) a routine taking all prognostic fields to a state that has evolved dynamically and that incorporates all vertical physics effects up to the turbulence scheme; (ii) a routine that executes the original turbulence, convection, and the cloud scheme on the grid columns not selected for SP; and (iii) a routine that executes all remaining physical processes after the cloud scheme that is being executed subsequently, for example, mass-fixers and diagnostics. We have also moved all stack-allocated arrays in the original horizontal loop to heap-allocated data structures, ensuring that the global model keeps its full state between these fractional time steps. The original time step is therefore equivalent to the consecutive execution of these three routines.

To disable the OpenIFS cloud and convection schemes as well as the boundary layer turbulence scheme for the superparameterized grid points, we have introduced a global *superparameterization mask*. All parameterization routines are being executed, but for grid points where the mask is set, the tendencies from these processes are set to zero so that their effects are discarded. In this approach additional diagnostics arising from these parameterizations can still be used, for example, the surface heat and momentum fluxes are




Figure 6. Zoom in on the superparameterized simulation and 3-D views on selected local models, compared to a satellite image from Terra/MODIS at 13 April 2012, 11:35 UTC. The star marks de Bilt, where radiosonde observations were performed.

still computed by the OpenIFS surface scheme and transferred to the local models to provide boundary conditions at the beginning of each SP time loop.

3.2. Interface to DALES

A similar library interface was created for DALES, with functions for initializing the model, setting tendencies, evolving the model until a given time, and fetching vertical profiles of the model variables. Creating a library interface for DALES required less involved changes than for OpenIFS. The necessary changes were mainly to add interface functions for tendencies and for retrieving horizontal averages of the variables.

3.3. OMUSE Coupling Framework

OMUSE (Pelupessy et al., 2017) is a framework for creating Python interfaces for scientific codes written in various languages such as C or Fortran. With OMUSE, we created Python versions of the function interfaces to OpenIFS and DALES described above. OMUSE is MPI-aware, making it possible to transparently communicate with MPI-parallelized models.

Through the OMUSE interface, both models can be controlled from a Python program by calling Python functions. Internally, OMUSE translates these function calls to Fortran function calls in the model codes, using MPI. Having MPI as the communication channel between the coupler and the models enable the models to be distributed over multiple nodes in a cluster. Furthermore, OMUSE hides the parallel nature of the models—every function in the OMUSE interface is collective over the MPI tasks of a given model, freeing the coupler from dealing with lower-level details of how the models are parallelized. Using OMUSE is a way to keep the model modular, making it relatively easy to, for example, substitute DALES with another LES code, or even a single-column cloud model.

3.4. The Coupling Code

The SP couplings described in section 2 have been implemented in a Python program using the OMUSE interfaces to OpenIFS and DALES. Figure 4 shows the interaction of the different components during a time step of the combined model.

Our setup does not require communication between the DALES instances and OpenIFS directly. All interactions are transferred through the coupling program, which fetches and compares the model states and provides feedback to the models in the form of tendencies. For localized SP the communication overhead

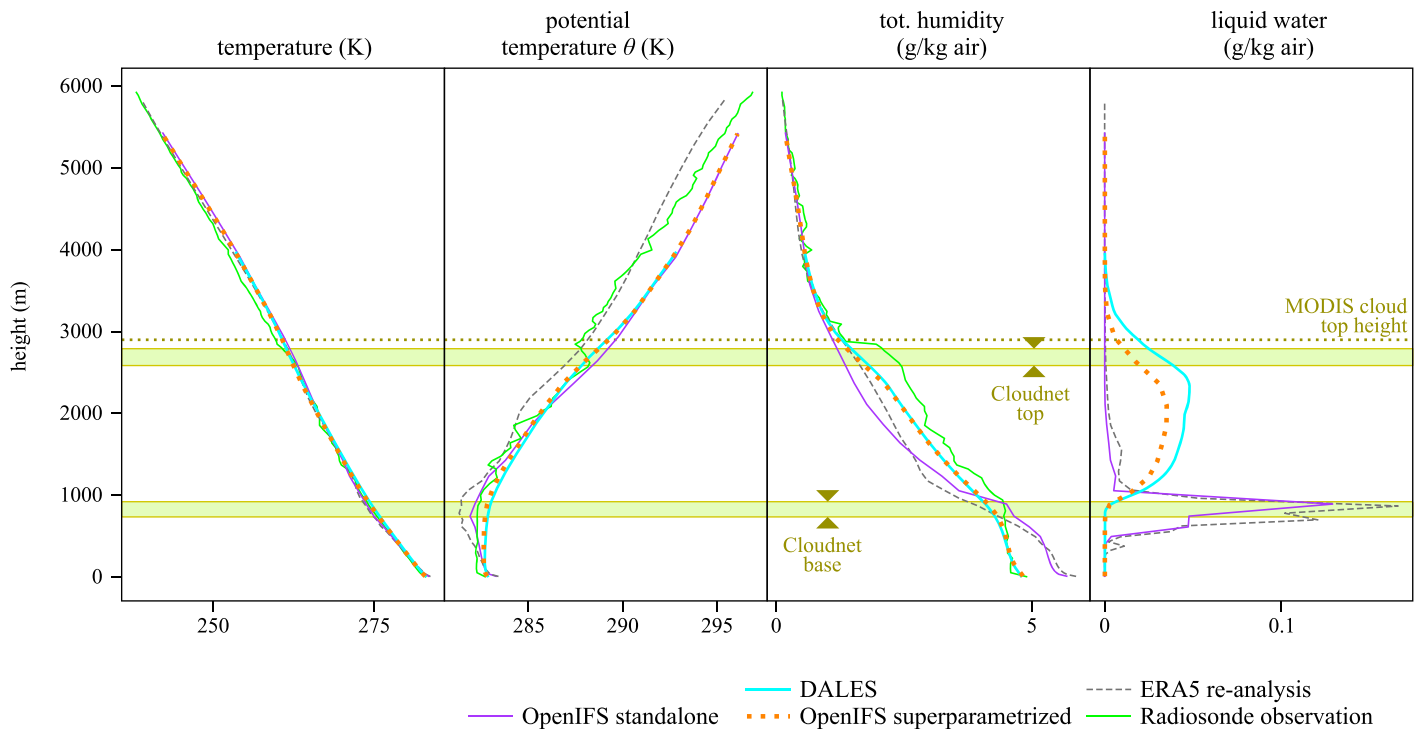


Figure 7. Vertical profiles at 13 April 2012, 12:00 UTC, of the grid point closest to de Bilt where radiosonde observations were performed. As a reference we use the ERA5 reanalysis and radiosonde observations. The horizontal line shows the cloud top height retrieved from MODIS. The wide horizontal lines show the range of cloud top and cloud base heights over a half-hour interval from Cloudnet, recorded by the cloud radar at the Cabauw site, 22 km Southwest of de Bilt.

remains limited. Given the way the SP coupling is formulated, no 3-D fields need to be exchanged—vertical profiles are sufficient. The cost of exchanging vertical profiles is generally small compared to the DALES run time. Having a separate coupling code in Python allows rapid prototyping and easy output of tailored diagnostics of the exchanged tendencies. It also keeps the code modular and easier to maintain.

The OMUSE framework supports (at least on many HPC clusters) dynamic instantiation of models; the coupling script is built to launch DALES instances within a user-defined area at initialization, making it easy to select an area for SP. Furthermore, unit conversions and vertical interpolation of profiles are implemented in the Python driver code. The system supports collection of basic performance statistics and adding a spin-up period at the start of the simulation, where the DALES instances are run for a specified time while being relaxed toward the vertical profiles of the OpenIFS model.

We note that precipitation within the local models is—for technical reasons—not yet coupled back to the global model's diagnostic rain fields. In the current implementation, the global model physics schemes that are affected by precipitation (the radiation and soil schemes) use the precipitation computed with the global model's parameterization of clouds and convection also within the superparameterized domain. The temperature and humidity tendencies from this parameterization scheme are switched off, and the tendencies are instead calculated in the local model.

Conversely, the precipitation generated within the local models act as a sink of humidity in the global model through the model coupling described in section 2.1, but the precipitation rates themselves are not used by other parameterizations in the global model. This inconsistency may be of importance for cases where soil properties and radiative fluxes are heavily influenced by precipitation.

4. Results

In this section we show initial results from the superparameterized model, in order to demonstrate the feasibility and potential of our approach and to evaluate the model's computational performance. We show




Figure 8. A time series of hourly vertical profiles of cloud liquid water content, at Cabauw (13 April 2012). The horizontal lines show the cloud top and cloud base heights, as reported by Cloudnet.

comparisons of the simulation results with observations and reanalysis data, while leaving a more extensive validation of the model for a later study.

Simulations using OpenIFS with a superparameterized setup over the Netherlands are performed for 13 April 2012. This day was characterized by a west-northwesterly flow steered by a low over the Northern part of the North Sea. Clear skies are observed over the relatively cold water of the North Sea, but convection developed when the cold air was advected from the North Sea over the Netherlands resulting in developing shallow cumulus clouds between 1 and 3 km over land. This case was selected because this type of boundary clouds is difficult to parameterize realistically and are not resolved by conventional CRM-based SP formulations. Thus, for this type of moist convection we expect a clear benefit from our approach. It is also interesting to explore to what extent the present SP formulation is capable of reproducing the larger-scale pattern with the cumulus cloud fields over the Netherlands and the clear boundary layers over most of the the North Sea interrupted by isolated cloud patches.

The OpenIFS initial state was constructed by interpolation from the reanalysis data set ERA-Interim (Dee et al., 2011). The local models were initialized with the vertical profiles of their corresponding grid points in the global model, with noise added in the horizontal direction to break the symmetry. After initialization, a spin-up of the local models was performed, where they were run for 4 hr while being relaxed toward the state of the global model. After spin-up the actual simulation was started, and the global and local models were time stepped together as described in section 2. OpenIFS was run with a T511L91 grid in all cases, giving a distance between neighboring grid points of about 40 km. The simulation parameters are summarized in Table 2. For the superparameterized run, we choose the extent of the DALES domains to match the grid point distance in OpenIFS.

4.1. Cloud Cover

We compare the appearance of the cloud fields from the simulation with satellite images and with results from the unmodified OpenIFS. Figure 5 shows a snapshot of the superparameterized simulation, compared




Figure 9. Maps of liquid water path and cloud cover, for the superparametrized model and stand-alone OpenIFS, compared with MODIS observations, 13 April 2012, 11:35 UTC. In the bottom row, low cloud cover is for the vertical range 0 ... 1,850 m, while medium cloud cover is for 1,850 ... 6,000 m approximately.

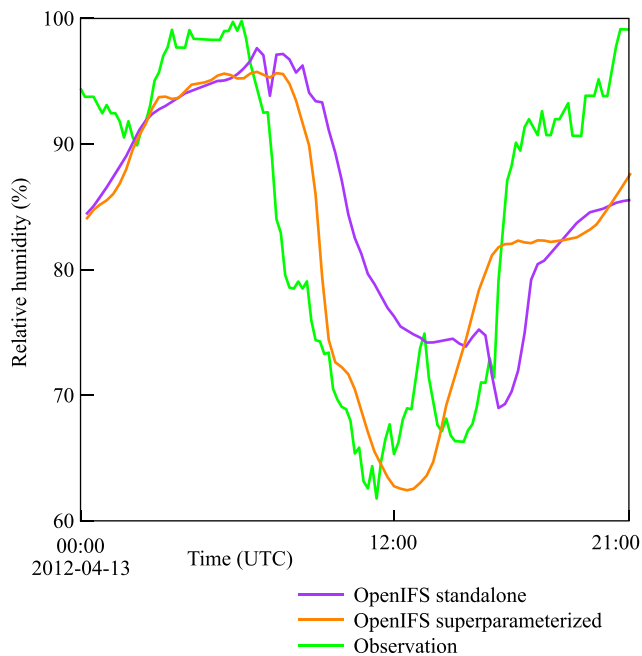


Figure 10. Time series of relative humidity at 2 m height compared with observations at the Cabauw site.

to a similar run of regular OpenIFS. The state of the simulations are shown at 11:35 UTC (this time was chosen to coincide with the overpass of the Terra satellite for comparison). The simulations were initialized at midnight UTC the same day. Animations of the simulations shown are available as supporting information.

The figures illustrate the clouds in OpenIFS and in DALES through the liquid water path. Since the cloud optical thickness has a nonlinear dependence on the liquid water path, we use a nonlinear color map for the DALES instances (a gamma correction with $\gamma = 1/2$, so the pixel coloring is determined by $(q_L/q_{L,max})^\gamma$). The plot still offers only a crude approximation of how the clouds would actually appear on a satellite image, since the cloud optical thickness depends strongly on the droplet size distribution, which is not taken into account in the plot. A more detailed quantitative comparison follows in section 4.2.

A magnification of the superparameterized run are shown in Figure 6 together with a satellite image of the same area. As can be seen in the figure, the local models over land show more clouds than the ones over the sea, in agreement with the satellite image and reanalysis data. However, large, stratiform cloud fields are underrepresented in the superparameterized OpenIFS.

When comparing DALES cloud fields with observations, it is good to keep in mind that the local models in a superparameterized setup give a representation of the convection and clouds at a grid point of the global model

but cannot be expected to accurately reproduce individual features or clouds seen in observations. One reason for this is that the initial state does not provide any small-scale information—the DALES simulations are initialized with vertical profiles from the global model. Second, the DALES simulations are performed with periodic boundary conditions, so that spatial coordinates in them do not directly correspond to any particular geographic coordinates.

4.2. Vertical Profiles

Vertical profiles of temperature, humidity, and liquid water humidity at a single OpenIFS grid point are shown in Figure 7. The superparameterized run is shown with profiles from both DALES and the corresponding grid point in OpenIFS. It can be seen that in the superparameterized run, DALES and OpenIFS are consistent with each other as can be expected from the coupling scheme. The superparameterized run is compared to a similar run of the nonsuperparameterized OpenIFS, to the reanalysis from ERA5, and to radiosonde observations. The largest difference in results is seen in the liquid water profiles (right panel), where the SP produces significantly higher clouds than standard OpenIFS. The SP result agrees well with the cloud top height measurement of MODIS, 2,900 m. Also, the total humidity measured by the radiosonde shows a sharp step at this height, consistent with this being the cloud top height. We additionally compare the liquid water result with cloud radar recordings from Cloudnet (Illingworth et al., 2007), taken at the Cabauw site, 22 km Southwest of de Bilt. Over 30 min, the cloud top height was measured at $2,690 \pm 100$ m, and the cloud base at 830 ± 90 m. These ranges are indicated with green stripes in Figure 7 and show a good agreement with the liquid water results from the superparameterized simulation. When comparing simulation results with reanalysis, one should remember that the reanalysis was done with IFS, of which OpenIFS is a version. There may thus be a bias for the reanalysis to behave similar to OpenIFS.

To further illustrate the difference between stand-alone OpenIFS and the superparameterized model, we plot a sequence of hourly cloud liquid water profiles at the Cabauw site in Figure 8. The observed values for the cloud top and cloud base heights, from Cloudnet, are shown as horizontal lines. The superparameterized run generally produces higher clouds and shows a better agreement with the observed cloud heights. Three additional runs were performed for other starting dates. Time series of liquid water profiles of 48 hr each are shown in the supporting information (Figures S1–S3). Also, in these runs there are cases where the models generate different amounts of vertical transport. Toward the end of Figure S1, SP generates higher clouds similar to the result in Figure 8. At the beginning of the run in Figure S1, and the end of the one in Figure




Figure 11. Timing diagram for the superparameterized run, showing three OpenIFS time steps around the time of the snapshot shown in Figure 5. The blue horizontal bars show the wall-clock time required for each individual DALES instance to complete the time step. The vertical bars show the OpenIFS computation (purple) and communication between the coupler and the models (orange).




Figure 12. Root-mean-square difference between the superparameterized runs and the ERA-5 reanalysis, over all superparameterized grid points for the vertical levels up to the height of the DALES models.

S3, the situation is reversed and stand-alone OpenIFS generates higher clouds than the superparameterized simulation.

Figure 9 shows maps of the liquid water path and the cloud cover for the superparameterized model, stand-alone OpenIFS, and observations from MODIS. Also differences between models and observation are shown. The dashed outline shows the superparameterized region. For the comparison, the MODIS data were coarse grained onto the same grid as OpenIFS, using Gaussian-weighted averaging with the full width at half maximum chosen as 40 km, corresponding to the grid spacing.

For the liquid water path, both models show similar deviations from the observation. The superparameterized run has less liquid water in general and shows a shift of clouds further in-land. Both models show less cloud cover than the observation, with SP showing less than stand-alone OpenIFS. The lowest two panels separately compare the cloud cover for low- and medium-height clouds between the two models and show that SP generally produced higher clouds in this case as was also seen in the vertical profiles in Figures 7 and 8

Figure 10 shows the relative humidity at 2 m height at the Cabauw site, and simulation results for the superparameterized model and stand-alone OpenIFS. At noon the superparameterized run shows a better agreement with observations, capturing the time and the depth of the drop in relative humidity. Thereafter, both models are roughly equally accurate.

While the comparison presented here is certainly too limited to draw broad conclusions about the accuracy of the superparameterized simulation, the match between the superparameterized clouds and MODIS, the radiosonde observations, and the Cloudnet cloud top and base heights

is encouraging. At the same time, the comparisons show that in particular the liquid water profiles from the reanalysis are inconsistent with the MODIS and Cloudnet observations. We note that the parameters of DALES have not been tuned for this comparison.

4.3. Performance

The superparameterized run with the 42 larger DALES models presented above, consisting of 21 hr of simulation and 4 hr of DALES spin-up, took 39 hr on 10 nodes of the Cray-XC40 system at ECMWF. Each node in this system contains two 18-core Intel Xeon EP E5-2695 V4 Broadwell processors. Each DALES instance used eight MPI tasks, resulting in 336 DALES processes in total. The first node ran the coupling script, OpenIFS with 17 processes, and 12 DALES processes. The remaining nine nodes ran 36 DALES processes each. The total cost of this simulation was 14,041 core-hours.

Figure 11 shows the distribution of computing time over OpenIFS, the DALES models, and the communication, around the time of the snapshot shown in Figure 5. Most of the time, 82% is spent on the DALES models, followed by communication and coupling with 16%. OpenIFS itself consumes only 2% of the total time. There is considerable variation in the computational times for the different DALES instances, which appear since DALES uses adaptive time stepping. The more convection there is in a DALES volume, the shorter the time step needs to be, and the longer the simulation takes. This results in some work imbalance, since for each time step of the global model, the whole calculation must wait until the slowest DALES instance completes. Dynamical balancing of the DALES work between SP time steps is in principle possible within our setup, but

Table 3
Comparison of the Computational Time for Different Parts of the Simulation, With and Without Acceleration Methods

LES grid	Time acc.	DALES	Run time (1,000 s)			Speedup factor	
			OpenIFS	Coupler	Total	DALES	Total
200 × 200 × 160	—	115.4	2.6	22.0	140.0	1.0	1.0
64 × 64 × 160	—	13.1	2.6	19.2	34.9	8.8	4.0
64 × 64 × 160	2	6.9	2.6	19.2	28.7	16.7	4.9

Note. All runs were performed as described in section 4.3, on the same computer system.

it will require the transfer of full DALES states across nodes, and its feasibility will be a subject of future investigation. Some further performance can most likely still be gained by carefully optimizing the job layout and also by overlapping some of the communication steps with computation. In the following section, we address ways of accelerating the local models themselves.

4.4. Acceleration

To reduce the computational cost of the superparameterized simulation, we consider ways of accelerating the local models, which consume most of the computational time. First, the horizontal extent of the local models may be chosen smaller than the global model's grid size (Xing et al., 2009, calls this the “reduced space strategy”). Global SP studies have also frequently used this strategy, sometimes combined with making the local models two-dimensional.

Second, the local models can be accelerated in time, using the mean state acceleration method of Jones et al. (2015), which was also used in the SP context by Parishani et al. (2017) with promising results. Briefly, this method assumes a separation of time scales, between the time of eddy motion (fast) and the time on which the local model's horizontal averages change (slow). The technique is a good match with SP, since only the horizontal averages, with the slow time scale, are coupled to the global model. In the mean state acceleration technique, after every time step in the local model, the horizontal averages of the tendencies are calculated. These average tendencies are then applied to the model variables in a horizontally uniform way, in order to accelerate the rate of change for the horizontal means.

These techniques are demonstrated with two accelerated SP runs with the same initial conditions as in the previous section. In the first, the horizontal extent of DALES is reduced to 64×64 columns, reducing the area to 10% of the original. In the second run, the mean state acceleration is also applied, with an acceleration factor of 2. To evaluate the accuracy of the accelerated runs, we plot the time evolution of the root-mean-square difference between the global model variables and the ERA-5 reanalysis over all superparameterized grid points (Figure 12). The plots show that the differences introduced by the acceleration are rather small.

The computational time requirements with and without acceleration are shown in Table 3. Reducing the area covered by DALES to 10% of the original reduces the DALES run time by a factor of 8.8. Adding mean state acceleration with an acceleration factor of 2 further reduces the DALES run time by a factor of 1.9. The total run time decreases less dramatically, by a factor 4 or 4.9, since the amounts of time spent on coupling and on OpenIFS remain constant. A simple way to decrease the time fraction spent on the coupling and communication is to allocate fewer processes per DALES—for this comparison we kept the job layout the same for all runs. Another interesting optimization possibility is to increase the concurrency of computation in the model codes with communication between the coupler and the models.

These acceleration results seem promising for reducing the computational cost of SP. We also note that the degrees of acceleration that can be achieved may depend on the case.

5. Conclusions and Outlook

We have demonstrated a SP of a global atmospheric model with a three-dimensional, high-resolution atmospheric LES over a configurable region. We show an example with a 21-hr run where a region with an area of 240 km × 280 km is superparameterized with high-resolution LES models. Reducing the extent of the local models and applying mean state acceleration drastically reduces the computational demands, with only minor deterioration of the results.

The coupling between the global and local models was implemented with a coupling program in Python, communicating with the model instances using the OMUSE framework. This approach makes the both SP algorithm and the allocation of computational resources explicit at the master script level. Controlling the allocation of resources to the different components of the coupled model is essential due to the computational imbalance between superparameterized and ordinary grid points. Moreover, the coupler script provides a convenient environment for exploratory work on the mapping of the exchanged tendencies and profiles between the global and local models.

Without acceleration methods, the major part of the computation is spent in time stepping DALES. With the introduction of acceleration methods for the local models, the performance has increased to the point where the time spent in the coupler becomes significant, giving us a motivation to address the coupler performance in the future. Another potential avenue for further improvement of the coupled model performance is the application of dynamical load-balancing strategies to the multiscale system, so one can benefit from the fact that some DALES instances finish earlier than others, due to their adaptive time stepping.

The simulation results demonstrate the potential of our approach. The cloudy regions, observed by MODIS, are reproduced by the superparameterized grid boxes. Comparison with the local observations at the Cabauw site shows that the superparameterized version of OpenIFS simulates a deeper convective mixing leading to improved profiles of especially the specific humidity and the cloud amount, compared with the standard version of the OpenIFS without SP.

The simulations also illustrate limitations of this approach. As can be observed from Figure 6, the superparameterized simulation poorly represents the observed coherent cloud structures of sizes comparable with or beyond the resolution of the global model (40 km). This is in part the consequence of the periodic boundary conditions of the local model and the lack of a direct coupling between the neighboring local model instances. This prevents the growth of mesoscale cloud structures, emerging from the smaller turbulent scales resolved in the local model, beyond the domain size of the local model. In addition the coupling between the local and the global model, as expressed by equations (2) and (4), is exclusively formulated in terms of the mean values of the prognostic (thermo)dynamic variables. Therefore, the coupling does not include any scale interaction of the variances of these variables, though many observational studies have shown a continuous growth of variances of temperature and humidity with the spatial scale as $\ell^{2/3}$, up to scales of several hundreds of kilometers without any scale break (Kahn et al., 2011). Work is therefore in progress to introduce an additional coupling of the variance of q_t between the local and global model. Other limitations are the coarser vertical resolution of the global model that might degrade sharp inversions resolved by local model and the radiation scheme that acts on the coarser grid of the global model. These latter limitations will likely have a degrading effect on the representation of stratocumulus-topped boundary layers but are expected to play a minor role in the present case.

As already mentioned in the introduction, this framework can serve as a useful benchmark for the development of new parameterization approaches. Many of the parameterization developments of cloud-related processes of the last 20 years have been guided by LES studies of relevant cases driven by realistic large-scale forcings (Brown et al., 2002; Siebesma et al., 2003; vanZanten et al., 2011). The present framework provides new opportunities in this respect. It provides realistic benchmarks over longer periods, over larger areas, with realistic forcings that are easy to set up. The framework also allows the use of different local models, for example, an alternative parameterization package, data-driven algorithms trained by the LES (Dorrestijn et al., 2013; Schneider et al., 2017), or conceptual mixed layer models (Caldwell et al., 2013). This way it is possible to test and compare different approaches which all are in balance with the large scales due to the interactive coupling.

By increasing the resolution of the global model and, accordingly, reducing the domain size of the local model, the present framework can also be used to quantify how the response of the local model will change. This will provide guidance for at which resolutions and for which processes scale-aware parameterizations are required. Such experiments will also be useful in exploring how mesoscale organization is emerging. By varying the resolution, the effect of imposing a scale break at different spatial scales on the mesoscale organization can be systematically explored.

Table A1
References for the different codes used in the superparameterization setup

Name	URL	Archived snapshot, DOI
sp-coupler	https://github.com/CloudResolvingClimateModeling/sp-coupler	10.5281/zenodo.1968305
OMUSE v1.1	https://bitbucket.org/omuse/omuse/	10.5281/zenodo.1407941
AMUSE	http://www.amusecode.org/	archived with OMUSE
DALES with OMUSE interface	https://github.com/CloudResolvingClimateModeling/dales	10.5281/zenodo.1345110
OpenIFS with OMUSE interface	https://git.ecmwf.int/scm/~g.vandenoord_esciencecenter.nl/oifs40r1-lib.git (requires license and a user account at ECMWF)	
OpenIFS, licensing information	https://confluence.ecmwf.int/display/OIFS	

Note. ECMWF = European Centre for Medium-Range Weather Forecasts.

Appendix A: Analysis of the Coupling Scheme

Here we analyze the model coupling scheme in more detail. We show that the desired equality (1) between the global and the (averaged) local model state is achieved in cases where the advection and source terms are entirely situated either in the local model or in the global model.

Consider Steps (i) to (iv) of the SP scheme given in section 2. Combining (i) and (ii) gives

$$Q(T + \Delta T) = \langle q(T) \rangle + \Delta T(A_Q(T) + S_Q(T)). \quad (\text{A1})$$

In (iv), $f_q(T)$ does not change with t , and thus, $\sum_{t=T}^{T+\Delta T} \Delta t f_q(T) = \Delta T f_q$. With (iii) and (iv), this leads to $q(T + \Delta T) = q(T) + Q(T + \Delta T) - \langle q(T) \rangle + \sum_{t=T}^{T+\Delta T} \Delta t(a_q(t) + s_q(t))$. Taking the horizontal average, we find

$$\langle q(T + \Delta T) \rangle = Q(T + \Delta T) + \sum_{t=T}^{T+\Delta T} \Delta t \langle a_q(t) + s_q(t) \rangle. \quad (\text{A2})$$

In general, $Q(T) \neq \langle q(T) \rangle$ and $Q(T + \Delta T) \neq \langle q(T) \rangle$, as can be seen from the identities just derived. Thus, the equality (1) is generally not satisfied. However, if all of the advection and sources with nonzero average are accounted for in one model (global or local) and none in the other, the equality is satisfied (albeit possibly in a time-lagged sense). More precisely, assume $A_Q(T) + S_Q(T) = 0$; that is, all advection and sources are in the local model. Then $Q(T + \Delta T) = \langle q(T) \rangle$ by construction. Conversely, if the local model has no advection or source terms with nonzero horizontal average, so that $\langle a_q(t) + s_q(t) \rangle = 0$ for all t , we have $\langle q(T + \Delta T) \rangle = Q(T + \Delta T)$.

If both $A_Q(T) + S_Q(T) \neq 0$ and $\langle a_q(t) + s_q(t) \rangle \neq 0$, we can consider the difference between $Q(T + \Delta T)$ on the one hand and a weighted average of $\langle q(T) \rangle$ and $\langle q(T + \Delta T) \rangle$ on the other hand. Defining the weighting parameter α with $0 \leq \alpha \leq 1$, we have

$$Q(T) - (\alpha \langle q(T) \rangle + (1 - \alpha) \langle q(T + \Delta T) \rangle) = \alpha \Delta T (A_Q(T) + S_Q(T)) - (1 - \alpha) \sum_{t=T}^{T+\Delta T} \Delta t \langle a_q(t) + s_q(t) \rangle. \quad (\text{A3})$$

The RHS equals zero if $\Delta T (A_Q(T) + S_Q(T))$ and $\sum_{t=T}^{T+\Delta T} \Delta t \langle a_q(t) + s_q(t) \rangle$ have the same sign, and their ratio equals $(1 - \alpha)/\alpha$. To satisfy the latter requirement, α must depend on time T , vertical level z , and prognostic variable q, Q .

Code Availability

All the codes required for the simulation are available under open-source licenses, except OpenIFS (for which a license can be requested from ECMWF). Table A1 lists the codes used, URLs for repositories, and DOI numbers for archived snapshots.

The top-level coupler code in Python is called `sp-coupler`. To run, it requires OMUSE and versions of OpenIFS and DALES, which include the OMUSE interfaces. For installation instructions, see the documentation in the coupler repository. The coupler repository includes a Singularity recipe. Singularity (Kurtzer et al., 2017) is a software container system for scientific applications. The Singularity recipe is used to build a singularity image—a self-contained unit containing all the programs and all their dependencies needed to run the simulation. However, building the Singularity image still requires access to OpenIFS source code from ECMWF.

Author Contributions

D. C. and P. S. conceived the project. F. J., G. v. d. O., D. C., and P. S. defined the coupling procedure. F. J., G. v. d. O., and I. P. wrote the coupling code and the OMUSE interfaces to DALES. G. v. d. O. created the OMUSE interface to OpenIFS. F. J. ran the simulations. J. H. G. and F. J. developed the visualizations. J. H. G. drew the figures. F. J. wrote the article text, with contributions and editing by all other authors.

Acknowledgments

We thank Glenn Carver at the ECMWF for help with OpenIFS and for providing us with initial states for the simulation. We also thank Maria Chertova for comments on the manuscript and Bram van Es for work on the setup and testing of the coupled model. Henk Klein Baltink at the Royal Netherlands Meteorological Institute KNMI assisted us with Cloudnet data. This work was supported by the Netherlands eScience Center (NLeSC) under Grant 027.015.G03. Furthermore, we acknowledge the use of ECMWF's computing and archive facilities in the research reported here. Also, SURFsara provided computing resources. The ERA5 and ERA-Interim data were provided by Copernicus Atmosphere Monitoring Service information 2018. Neither the European Commission nor ECMWF is responsible for any use that may be made of the Copernicus information or data it contains. We acknowledge the use of imagery from the NASA Worldview application (<https://worldview.earthdata.nasa.gov/>) operated by the NASA/Goddard Space Flight Center Earth Science Data and Information System (ESDIS) project. The Terra/MODIS cloud data was obtained from the Level 1 and Atmosphere Archive and Distribution System (LAADS) Distributed Active Archive Center (DAAC), located in the Goddard Space Flight Center in Greenbelt, Maryland (<https://ladsweb.nascom.nasa.gov/>). Radiosonde data were provided by KNMI. It was recorded with a Vaisala RS92 radiosonde, released from de Bilt, the Netherlands. The maps were drawn with Matplotlib (Hunter, 2007) with Basemap and Cartopy. The 3-D cloud images were raytraced in Blender (<https://www.blender.org/>).

References

Benedict, J. J., & Randall, D. A. (2009). Structure of the Madden-Julian oscillation in the superparameterized CAM. *Journal of the Atmospheric Sciences*, *66*(11), 3277–3296. <https://doi.org/10.1175/2009JAS3030.1>

Böing, S. J., Jonker, H. J. J., Siebesma, A. P., & Grabowski, W. W. (2012). Influence of the subcloud layer on the development of a deep convective ensemble. *Journal of the Atmospheric Sciences*, *69*(9), 2682–2698. <https://doi.org/10.1175/JAS-D-11-0317.1>

Bony, S., & Dufresne, J.-L. (2005). Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models. *Geophysical Research Letters*, *32*, L20806. <https://doi.org/10.1029/2005GL023851>

Bony, S., Stevens, B., Frierson, D. M. W., Jakob, C., Kageyama, M., Pincus, R., et al. (2015). Clouds, circulation and climate sensitivity. *Nature Geoscience*, *8*, 261. <https://doi.org/10.1038/ngeo2398>

Bretherton, C. S., & Khairoutdinov, M. F. (2015). Convective self-aggregation feedbacks in near-global cloud-resolving simulations of an aquaplanet. *Journal of Advances in Modeling Earth Systems*, *7*, 1765–1787. <https://doi.org/10.1002/2015MS000499>

Brown, A. R., Cederwall, R. T., Chlond, A., Duynkerke, P. G., Golaz, J.-C., Khairoutdinov, M., et al. (2002). Large-eddy simulation of the diurnal cycle of shallow cumulus convection over land. *Quarterly Journal of the Royal Meteorological Society*, *128*(582), 1075–1093. <https://doi.org/10.1256/003590002320373210>

Caldwell, P. M., Zhang, Y., & Klein, S. A. (2013). CMIP3 subtropical stratocumulus cloud feedback interpreted through a mixed-layer model. *Journal of Climate*, *26*(5), 1607–1625. <https://doi.org/10.1175/JCLI-D-12-00188.1>

Carver, G., & Vana, F. (2017). OpenIFS home. <https://confluence.ecmwf.int/display/OIFS/>

Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., et al. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, *137*(656), 553–597. <https://doi.org/10.1002/qj.828>

Dorrestijn, J., Crommelin, D. T., Siebesma, A. P., & Jonker, H. J. (2013). Stochastic parameterization of shallow cumulus convection estimated from high-resolution model data. *Theoretical and Computational Fluid Dynamics*, *27*(1-2), 133–148. <https://doi.org/10.1007/s00162-012-0281-y>

ECMWF (2014a). IFS documentation. Part III: Dynamics and numerical procedures, CY40R1. Retrieved from <https://www.ecmwf.int/sites/default/files/elibrary/2014/9203-part-iii-dynamics-and-numerical-procedures.pdf>

ECMWF (2014b). IFS documentation. Part IV: Physical processes, CY40R1. Retrieved from <https://www.ecmwf.int/sites/default/files/elibrary/2014/9204-part-iv-physical-processes.pdf>

Grabowski, W. W. (2001). Coupling cloud processes with the large-scale dynamics using the cloud-resolving convection parameterization (CRCP). *Journal of the Atmospheric Sciences*, *58*(9), 978–997. [https://doi.org/10.1175/1520-0469\(2001\)058h0978:CCPWTLi2.0.CO;2](https://doi.org/10.1175/1520-0469(2001)058h0978:CCPWTLi2.0.CO;2)

Grabowski, W. W. (2004). An improved framework for superparameterization. *Journal of the Atmospheric Sciences*, *61*(15), 1940–1952. [https://doi.org/10.1175/1520-0469\(2004\)061h1940:AIFFSi2.0.CO;2](https://doi.org/10.1175/1520-0469(2004)061h1940:AIFFSi2.0.CO;2)

Grabowski, W. W. (2016). Towards global large eddy simulation: Super-parameterization revisited. *Journal of the Meteorological Society of Japan. Ser. II*, *94*(4), 327–344. <https://doi.org/10.2151/jmsj.2016-017>

Grabowski, W. W., & Smolarkiewicz, P. K. (1999). CRCP: A cloud resolving convection parameterization for modeling the tropical convecting atmosphere. *Physica D: Nonlinear Phenomena*, *133*(1-4), 171–178. [https://doi.org/10.1016/S0167-2789\(99\)00104-9](https://doi.org/10.1016/S0167-2789(99)00104-9)

Heinze, R., Dipankar, A., Henken, C. C., Moseley, C., Sourdeval, O., Trömel, S., et al. (2017). Large-eddy simulations over Germany using ICON: A comprehensive evaluation. *Quarterly Journal of the Royal Meteorological Society*, *143*(702), 69–100. <https://doi.org/10.1002/qj.2947>

Heus, T., van Heerwaarden, C. C., Jonker, H. J. J., Pier Siebesma, A., Axelsen, S., van den Dries, K., et al. (2010). Formulation of the Dutch Atmospheric Large-Eddy Simulation (DALES) and overview of its applications. *Geoscientific Model Development*, *3*(2), 415–444. <https://doi.org/10.5194/gmd-3-415-2010>

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing In Science & Engineering*, *9*(3), 90–95. <https://doi.org/10.1109/MCSE.2007.55>

Illingworth, A. J., Hogan, R. J., O'Connor, E. J., Bouniol, D., Brooks, M. E., Delanoë, J., et al. (2007). Cloudnet: Continuous evaluation of cloud profiles in seven operational models using ground-based observations. *Bulletin of the American Meteorological Society*, *88*(6), 883–898. <https://doi.org/10.1175/BAMS-88-6-883>

Jones, C. R., Bretherton, C. S., & Pritchard, M. S. (2015). Mean-state acceleration of cloud-resolving models and large eddy simulations. *Journal of Advances in Modeling Earth Systems*, *7*, 1643–1660. <https://doi.org/10.1002/2015MS000488>

Jung, J. H., & Arakawa, A. (2010). Development of a Quasi-3d multiscale modeling framework: Motivation, basic algorithm and preliminary results. *Journal of Advances in Modeling Earth Systems*, *2*, 11. <https://doi.org/10.3894/JAMES.2010.2.11>

Kahn, B. H., Teixeira, J., Fetzer, E. J., Gettelman, A., Hristova-Veleva, S. M., Huang, X., et al. (2011). Temperature and water vapor variance scaling in global models: Comparisons to satellite and aircraft data. *Journal of the Atmospheric Sciences*, *68*(9), 2156–2168. <https://doi.org/10.1175/2011JAS3737.1>

Khairoutdinov, M., & Randall, D. A. (2001). A cloud resolving model as a cloud parameterization in the NCAR community climate system model: Preliminary results. *Geophysical Research Letters*, *28*(18), 3617–3620. <https://doi.org/10.1029/2001GL013552>

Khairoutdinov, M., Randall, D., & DeMott, C. (2005). Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes. *Journal of the Atmospheric Sciences*, *62*(7), 2136–2154. <https://doi.org/10.1175/JAS3453.1>

Kurtzer, G. M., Sochat, V., & Bauer, M. W. (2017). Singularity: Scientific containers for mobility of compute. *PLOS ONE*, *12*(5), 1–20. <https://doi.org/10.1371/journal.pone.0177459>

Marchand, R., & Ackerman, T. (2011). A cloud-resolving model with an adaptive vertical grid for boundary layer clouds. *Journal of the Atmospheric Sciences*, *68*(5), 1058–1074. <https://doi.org/10.1175/2010JAS3638.1>

Miyamoto, Y., Kajikawa, Y., Yoshida, R., Yamaura, T., Yashiro, H., & Tomita, H. (2013). Deep moist atmospheric convection in a subkilometer global simulation. *Geophysical Research Letters*, *40*, 4922–4926. <https://doi.org/10.1002/grl.50944>

Neggers, Roel A. J., Heus, T., & Siebesma, A. P. (2011). Overlap statistics of cumulusiform boundary-layer cloud fields in large-eddy simulations. *Journal of Geophysical Research*, *116*, D21202. <https://doi.org/10.1029/2011JD015650>

Parishani, H., Pritchard, M. S., Bretherton, C. S., Wyant, M. C., & Khairoutdinov, M. (2017). Toward low-cloud-permitting cloud superparameterization with explicit boundary layer turbulence. *Journal of Advances in Modeling Earth Systems*, *9*, 1542–1571. <https://doi.org/10.1002/2017MS000968>

Pelupessy, F. I., van Elteren, A., de Vries, N., McMillan, S. L. W., Drost, N., & Portegies Zwart, S. F. (2013). The astrophysical multipurpose software environment. *Astronomy & Astrophysics*, *557*, A84.

- Pelupessy, I., van Werkhoven, B., van Elteren, A., Viebahn, J., Candy, A., Portegies Zwart, S., & Dijkstra, H. (2017). The oceanographic multipurpose software environment (OMUSE v1.0). *Geoscientific Model Development*, *10*(8), 3167–3187. <https://doi.org/10.5194/gmd-10-3167-2017>
- Portegies Zwart, S., McMillan, S., Harfst, S., Groen, D., Fujii, M., Nualláin, Ó., et al. (2009). A multiphysics and multiscale software environment for modeling astrophysical systems. *New Astronomy*, *14*(4), 369–378. <https://doi.org/10.1016/j.newast.2008.10.006>
- Portegies Zwart, S. F., McMillan, S. L. W., van Elteren, A., Pelupessy, F. I., & de Vries, N. (2013). Multi-physics simulations using a hierarchical interchangeable software interface. *Computer Physics Communications*, *184*(3), 456–468. <https://doi.org/10.1016/j.cpc.2012.09.024>
- Schalkwijk, J., Jonker, H. J. J., Siebesma, A. P., & Van Meijgaard, E. (2015). Weather forecasting using GPU-based large-eddy simulations. *Bulletin of the American Meteorological Society*, *96*(5), 715–723. <https://doi.org/10.1175/BAMS-D-14-00114.1>
- Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth system modeling 2.0: A blueprint for models that learn from observations and targeted high-resolution simulations. *Geophysical Research Letters*, *44*, 12,396–12,417. <https://doi.org/10.1002/2017GL076101>
- Schneider, T., Teixeira, J., Bretherton, C. S., Brient, F., Pressel, K. G., Schär, C., & Siebesma, A. Pier (2017). Climate goals and computing the future of clouds. *Nature Climate Change*, *7*, 3. <https://doi.org/10.1038/nclimate3190>
- Siebesma, A. P., Bretherton, C. S., Brown, A., Chlond, A., Cuxart, J., Duynkerke, P. G., et al. (2003). A large eddy simulation intercomparison study of shallow cumulus convection. *Journal of the Atmospheric Sciences*, *60*(10), 1201–1219. [https://doi.org/10.1175/1520-0469\(2003\)60h1201:ALESISi2.0.CO;2](https://doi.org/10.1175/1520-0469(2003)60h1201:ALESISi2.0.CO;2)
- Tao, W. K., & Chern, J. D. (2017). The impact of simulated mesoscale convective systems on global precipitation: A multiscale modeling study. *Journal of Advances in Modeling Earth Systems*, *9*, 790–809. <https://doi.org/10.1002/2016MS000836>
- vanZanten, M. C., Stevens, B., Nuijens, L., Siebesma, A. P., Ackerman, A. S., Burnet, F., et al. (2011). Controls on precipitation and cloudiness in simulations of trade-wind cumulus as observed during RICO. *Journal of Advances in Modeling Earth Systems*, *3*, 19. <https://doi.org/10.1029/2011MS000056>
- Wicker, L. J., & Skamarock, W. C. (2002). Time-splitting methods for elastic models using forward time schemes. *Monthly Weather Review*, *130*(8), 2088–2097. [https://doi.org/10.1175/1520-0493\(2002\)130h2088:TSMFEMi2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130h2088:TSMFEMi2.0.CO;2)
- Xing, Y., Majda, A. J., & Grabowski, W. W. (2009). New efficient sparse space-time algorithms for superparameterization on mesoscales. *Monthly Weather Review*, *137*(12), 4307–4324. <https://doi.org/10.1175/2009MWR2858.1>