Weather Forecasting Using GPU-Based Large-Eddy Simulations

BY JERÔME SCHALKWIJK, HARMEN J. J. JONKER, A. PIER SIEBESMA, AND ERIK VAN MEIJGAARD

A HISTORY OF GLOBAL NUMERICAL WEATHER PREDICTION. The spectacular development of computational resources in the past decades has had a profound impact on the field of numerical weather and climate modeling. It has facilitated significant improvements in the description of key physical processes such as radiative transfer and has led to more accurate flow solvers. In addition, it has enabled sophisticated data-assimilation and ensemble-prediction schemes, both of which have turned out to be vital for improved prediction skill. Parallel to these developments, the increased computational power has led to a gradual but steady refinement of the computational grid. This has allowed models to resolve an increasingly large portion of the atmospheric scales of motion visualized in Fig. 1. The unresolved scales need to be approximated in a statistical way through statistical parameterizations, inevitably involving uncertain assumptions. The parametric representation of unresolved processes, especially processes related to clouds and convection, is considered to form a major source of uncertainty in weather (Slingo and Palmer 2011) and climate models (Dufresne and Bony 2008).

The historic evolution of computational grids is illustrated in the top panel of Fig. 2, which shows how the spatial scales treated by operational global numerical weather prediction (NWP) models have evolved in time. The range of resolved scales is visualized by a horizontal bar, with the largest scale (domain size) at the right and the smallest scale (resolution) at the left. The width of the bar is therefore a key measure of computational cost. Due to ever-increasing computational resources, operational NWP models have undergone an exponential increase in horizontal resolution. This growth started in 1974 with the model of the National Meteorological Center (NMC)
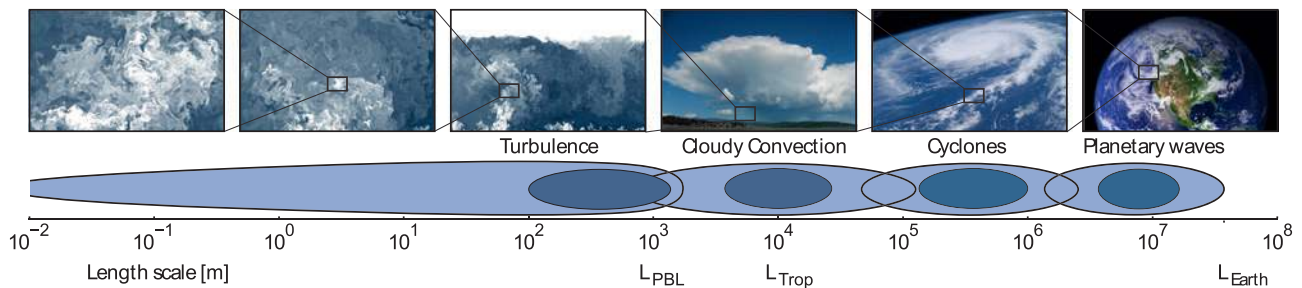

FIG. 1. Illustration of the atmospheric scales of motion, following Smagorinsky (1974). Common atmospheric phenomena are categorized according to their typical length scale (blue shades roughly indicate where the phenomena's energy is most often concentrated). Important length scales are L_{PBL} and L_{Trop}, the typical height of the planetary boundary layer and that of the troposphere, respectively.

AFFILIATIONS: SCHALKWIJK AND JONKER—Delft University of Technology, Delft, Netherlands; SIEBESMA—Royal Netherlands Meteorological Institute, De Bilt, and Delft University of Technology, Delft, Netherlands; VAN MEIJGAARD—Royal Netherlands Meteorological Institute, De Bilt, Netherlands

CORRESPONDING AUTHOR: Jerôme Schalkwijk, Department of Geoscience and Remote Sensing, Delft University of Technology, Stevinweg 1, 2628 CN Delft, Netherlands

E-mail: J.Schalkwijk@tudelft.nl
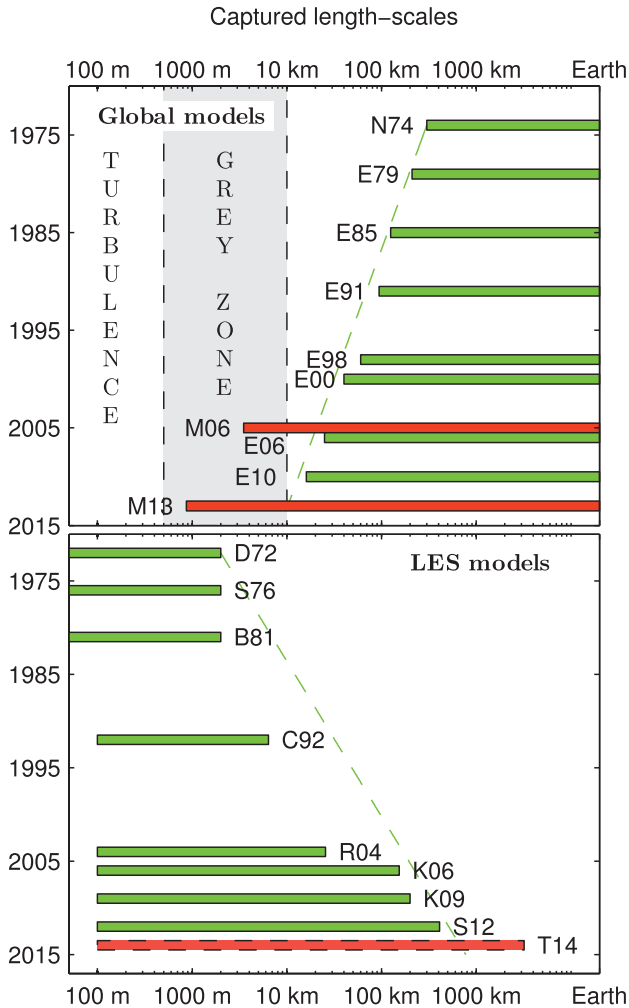
Captured length–scales

FIG. 2. Evolution of the scale range captured by numerical models. The top panel depicts global numerical weather prediction models, with green bars for runs in an operational setting and red bars for extraordinary simulations. The bottom panel depicts LES models, where the dashed-red bar shows the virtual possibilities of the Titan supercomputer. The tags refer to citations in the text.

at 300-km resolution (denoted N74; see Shuman 1989) and continued up to the resolution of 16 km that is now used by the latest version of the European Centre for Medium-Range Weather Forecasts model (E79–E10; see e.g., Simmons et al. 1989; European Centre for Medium-Range Weather Forecasts 2014). The red bars illustrate the computational breakthroughs by Miura et al. (M06; 2007), who simulated the global weather for one week at 3.5-km resolution, and by Miyamoto et al. (M13; 2013), who simulated 12 h at 0.87-km resolution some years later. While such exceptional cases cannot be performed on a

regular basis, they illustrate the very limit of what is currently possible.

Operational global NWP models are presently on the verge of using resolutions finer than the depth of the troposphere, L_{Trop} (10 km) (see Fig. 1). This implies that they are beginning to resolve the vertical convective overturning by cumulus clouds, but still need its partial parameterized representation. This obstacle, known as the "gray zone" or "Terra Incognita" (Wyngaard 2004) is like the proverbial "chasm" that cannot be crossed in small steps. Ideally the representation of convective overturning at these resolutions should be distributed smoothly (i.e., as a function of resolution), between the subgrid parameterizations on the one hand and explicit simulation on the other (Molinari and Dudek 1992; Wyngaard 2004; Arakawa et al. 2011). This can be achieved by making parameterizations "scale aware," but a general framework for such an approach is presently lacking within the context of NWP models (e.g., particularly for convection).

If resolutions of 100 m are used, the vertical overturning of cumulus convection is resolved, but one could argue that a next gray zone of three-dimensional turbulence is entered: whereas the largest turbulent eddies, having the size of the depth of the planetary boundary layer L_{PBL} (1 km), are resolved, the smaller turbulent eddies are still unresolved. Fortunately, the self-similar nature of inertial-range turbulence (with scales of 100 m or less), visualized in Fig. 1, can be used to express the transport of unresolved eddies in terms of the resolved eddies in a truly scale-aware fashion (see next section). It would therefore be a cornerstone achievement to perform global NWP at a turbulence-resolving grid resolution of 100 m. A naive extrapolation of the historical development displayed in the top panel of Fig. 2 suggests that this could be feasible on a daily basis around 2060.

DEVELOPMENT OF LES MODELS. Pioneered by Lilly (1962) and Deardorff (1972), Large-Eddy Simulation (LES) models employ subgrid transport parameterizations that exploit the self-similarity of inertial-range turbulence to represent atmospheric turbulence, instead of resolving motions down to the millimeter scale. Therefore, resolutions on the order of 100 m can suffice to simulate turbulent transport in the absence of external complexities, like detailed terrain features, although higher resolutions and therefore higher effective Reynolds numbers are desirable (see, e.g., Sullivan and Patton 2011). The historic development of computational grids for this type of model is shown

in the bottom panel of Fig. 2. Initially, efforts were concentrated on improving the model characteristics and including more model components (D72, S76, B81; Deardorff 1972; Sommeria 1976; Bougeault 1981), but since the mid-nineties, there has been interest in increasing the domain size (C92, R04; Chlond 1992; de Roode et al. 2004). An important subsequent development were the simulations by Khairoutdinov and Randall (K06; 2006) and Khairoutdinov et al. (K09; 2009), who extended the domain size beyond 100 km. In that respect, it is tempting to also extrapolate the trend in the lower panel and imagine when the scale of the Earth will be reached. Interestingly, such thought experiments demonstrate that the final result of both approaches is the same: a global turbulence-resolving model. The refinement of global models might seem the most natural approach to this challenge, as these models have all of the important components already in place, having been running in an operational setting for decades. The gray zone is not easily crossed, however, especially for the traditional hydrostatic models. Promising work is currently being undertaken to develop nonhydrostatic global models that resolve convection when the employed resolution and parameterizations allow it (e.g., the MPAS-model, Skamarock et al. 2012).

An alternative approach to the gray-zone problem may be to use LES modeling as a starting point and develop a framework to enlarge the employed domain. The idea of using such an LES-based enlargement approach was clearly expressed in a recently funded massive German initiative called HD(CP)2 (Rauser 2014), which formulated the aim to simulate the weather over Germany with a resolution of 100 m. However, the realization of such an ambition is a major challenge.

While focusing on the domain expansion of "classical" LES models and the grid refinement of NWP models, we entirely neglected an important body of intermediate-scale models (i.e., limited-area models, or mesoscale models). Mesoscale models (such as WRF) have

been developed that close the gap between NWP and LES models (e.g., Skamarock et al. 2008; Mirocha et al. 2013), or allow detailed representation of a specific subdomain through a nested LES (Moeng et al. 2007). Scale-aware turbulence models as developed by, for example, Bou-Zeid et al. (2005), Hatlee and Wyngaard (2007), and Perot and Gadebusch (2009) may allow such models to flexibly scale toward typical LES resolutions. For the sake of simplicity, however, we focus here on the two "ends" of the spectrum, highlighting the contrast between domain size and resolution.

WEATHER SIMULATION WITH LES? Here we perform an exploratory exercise into the LES-enlargement approach by employing the newest computer architectures—particularly graphical processing units (GPUs). GPUs have gone through tremendous development in the past decade, aided by the large financial stimulus from commercial gaming. In contrast to the central processing unit (CPU), GPUs were explicitly designed for parallel computing. Recently, GPUs evolved into general-purpose devices for massively parallel computations. In this form, they can provide an enormous computational boost, provided the programs are specifically adapted to the GPU design. A short explanation of GPU modeling is given in the sidebar.

Because relatively simple parameterizations facilitate quick adaptation to a new computer architecture, we were able to create GALES, a GPU-resident Atmospheric LES (Schalkwijk et al. 2012; Heus et al. 2010). As a result, simulations that normally require roughly 50 processors can be performed on a single GPU with comparable speeds. Yet, to enable the massive simula-


Fɪɢ. 3. The simulated cloud field by GALES on (a) 28 May 2006 (cloud streets) and (b) 28 Jun 2009 (thunderstorm).

tions presented below, GALES was further extended to exploit the multi-GPU hardware as featured by the latest "accelerated" supercomputers. The hardware in this study was provided by the French CURIE Hybrid supercomputer, allowing for simulations using as many as 256 GPUs simultaneously. This enabled us to simulate a domain that spans 400×400 km^2 with 100-m resolution, enough to cover the Netherlands.

To test the applicability of LES modeling to the daily weather, we have simulated three different "archetype" weather situations by nesting GALES in a weather model, as detailed in the appendix. These situations comprise fair weather cumulus clouds, cloud streets in high-wind conditions (Fig. 3a), and the development

FIG. 4. Clouds over the Netherlands on 6 Jul 2004. The image shows the simulated cloud field by GALES, with the corresponding satellite image as inset.

GPU COMPUTATION

The differences between the graphical processing unit (GPU) and the central processing unit (CPU) stem from their traditional roles in the computer. The CPU (the traditional "brains" of the computer) is responsible for the operating system (OS) and running programs. As a result, the CPU chip acquired a large cache (short-term, very fast memory for data in use) and a large scheduler

FIG. SB1. Computational scaling properties of GALES on multi-GPU hardware in terms of observed speed-up (with respect to using 1 GPU) after increasing the number of GPUs. The pluses (+) represent an extrapolation to the Oak Ridge GPU supercomputer, Titan, and the position of a global domain in this scaling diagram.

(responsible for distributing the processing power over the various tasks and programs) to efficiently handle a relatively small number of unrelated and very diverse tasks. In contrast, GPUs typically perform a large number of nearly identical computations—they are responsible for the intensive vector/matrix computations associated with computer graphics—and therefore developed into massively parallel devices.

LES models perform calculations that are very similar to these matrix computations. First, the equation solved is the same for every grid node, and second, the data are ordered in a structured manner. This allows for a straightforward task division on the one hand, and efficient memory access on the other hand: each consecutive GPU core performs the same instruction, but on a consecutive data element.

GALES is written in C++ and CUDA: CPU instructions in C++ for program management and I/O tasks, CUDA for the GPU computations. The implementation for single-GPU purposes is described in Schalkwijk et al. (2012). To perform the large simulations shown in this manuscript, we have complemented the GPU parallelism with the more traditional method of message passing. The domain is divided into subdomains, which are distributed among the processors (CPUs). Each CPU offloads its subdomain to the GPU, which performs computations for

Fig. 5. The simulated cloud field that was shown in Fig. 4, but rendered from a different perspective.

of a severe thunderstorm (Fig. 3b). For brevity, we will focus on the case of fair weather cumulus clouds (Fig. 4), which are the most abundant clouds on the globe (Rossow et al. 2005), but which are still poorly represented in conventional weather and climate models. Note that the simulations in this work require 4 h of computation for every simulated hour, and are therefore yet unfit as forecasts. Rather, they represent special cases of extreme computing and serve to illustrate what will become possible in the near future.

The cumulus case visualized in Fig. 4 is characterized by cumulus clouds that are typically 100 m–1 km wide and up to a few kilometers deep. A comparison with the satellite image at the same time (see inset in Fig. 4) shows that the simulation indeed forecasts the observed fair weather cumulus, which are notoriously

the subdomain. Communication between the subdomains is performed via the CPU, using Message Passing Interface (MPI), after every time-step. We follow the parallelization methodology described in Jonker et al. (2013), except that the 2D domain decomposition is performed over the horizontal directions (i.e., each GPU has the full vertical range) to efficiently calculate precipitation, and later, radiation.

Figure SB1 shows the computational "weak scaling" chart of GALES. It shows how the speed-up of the program improves with additional processor cores. Speed-up is a measure of the relative advantage of additional processors; "perfect scaling" implies full utilization of the additional computational power. In reality, communication between processor cores may prevent perfect scaling. As revealed by Fig. SB1, GALES adheres to the perfect scaling curve for over two orders of magnitude in the number of GPUs. This indicates that the simulations in the current study have not yet started saturating.

Since all cores (> 500!) within a GPU can access the GPU's memory, no data segmentation is required on the GPU level. Therefore, the number of data "blocks" that must be communicated between processors is much smaller than that for an equivalent simulation using CPUs only. This significantly reduces communication overhead. Later tests showed that when simulating $4{,}096^2 \times 256$ grid nodes using 256 GPUs, GALES spends

roughly 30% of the time on MPI transfer, and 10% on CPU-GPU transfer.

If all GPUs of Oak Ridge's supercomputer Titan (>16,384) could be used, a simulation of a $3{,}200 \times 3{,}200$-km^2 domain at 100-m resolution would already be possible today, suggesting that a global turbulence-resolving simulation could be possible in fewer than 10 years. Such simulations will not yet be *forecasts*: the simulations in this work typically required 4 h of computing to simulate 1 h, using 1.5-s time-steps. In the case of perfect weak scaling, this time-to-solution ratio remains constant; forecasting will thus require another ten-to-hundredfold increase in computational power (which still only amounts to 4–7 additional years).

Of course, many other challenges (e.g., complex terrain, pressure solver, data assimilation, etc.) also entail significant additional computational cost for LES models as they approach global simulations. Nevertheless, LES models have a number of key numerical advantages. First, their structure is very well suited for massively parallel systems. Second, whereas the equations in LES models are already integrated using time-steps of 1–10 s to resolve turbulent motions, operational weather models now feature time steps of roughly 15 min. For the latter models to become turbulence resolving, therefore, not only must the number of grid points be increased to reach 100-m resolution, but the size of their time-steps must decrease by a factor of roughly 400.

difficult to represent in conventional models (Slingo and Palmer 2011). The high-resolution features are illustrated yet more clearly in Fig. 5, which shows a three-dimensional rendering of the same cloud field. The shadows underneath the clouds illustrate the possibility of representing the fine interaction between clouds, radiation, and the surface. This interaction directly affects the temperature and humidity in the atmospheric boundary layer, and thereby the daily weather.

In addition, the sea breeze effects, recognizable by clear skies in the coastal regions and intensified cloudiness further inland, are remarkably well represented. However, the present study should by no means be considered a comprehensive skill analysis. Our purpose is to present a *proof of concept* of predictive Large-Eddy Simulations and to illustrate its potential for operational forecasting.

In general, it is well accepted that many of the long-term biases in weather and climate models are due to an incorrect representation of the interactions between clouds and the large-scale circulation (Stevens and Bony 2013). Such biases include the wrong phasing of precipitation in the diurnal cycle (Betts and Jakob 2002), the misrepresentation of the Madden–Julian oscillation (Miura et al. 2007), and systematic errors in the precipitation patterns in the tropics. These biases are likely due to an incorrect interaction between the large-scale resolved dynamics and the parameterized cloud processes in present-day weather and climate models. Case studies on smaller domains have shown that many of these biases can be mitigated by using high-resolution simulations.

CONCLUSION. Bearing in mind our earlier extrapolation anticipating high-resolution operational NWP models in the year 2060, it is useful to estimate when a global 100-m simulation will become technically possible in a special case of extreme computing. For GALES we know exactly how additional GPUs improve the simulation speed and the domain size (see sidebar). This leads us to conclude that the world's top GPU-accelerated supercomputer, Oak Ridge's Titan, could run a simulation of 3,200 x 3,200 km^2 domain at 100-m resolution already today (for reference, the Earth's diameter is roughly 6,400 km). This virtual breakthrough is illustrated with the red bar in the bottom panel of Fig. 2 (tagged T14). If, in addition, we extrapolate the historic supercomputing trend (Strohmaier et al. 2005), we find that a global turbulence-resolving simulation will be possible in less than 10 years.

Of course, this approach also faces challenges in its route to global simulations. In many cases, LES models require modifications to the pressure solver and even to the grid itself to handle the Earth's curvature and orography, not to mention the assimilation of observational data. Although such challenges might entail significant additional computational cost, they do not conflict with the model design: many issues have already been solved in the fields of weather forecasting or, for example, fluid dynamics. The relatively low number of assumptions in LES models aids the quick adaptation to new computational strategies, and the models are well prepared to utilize the exponential increase of computational power. Therefore, we should not be surprised when the first global LES experiments appear within the next decade.
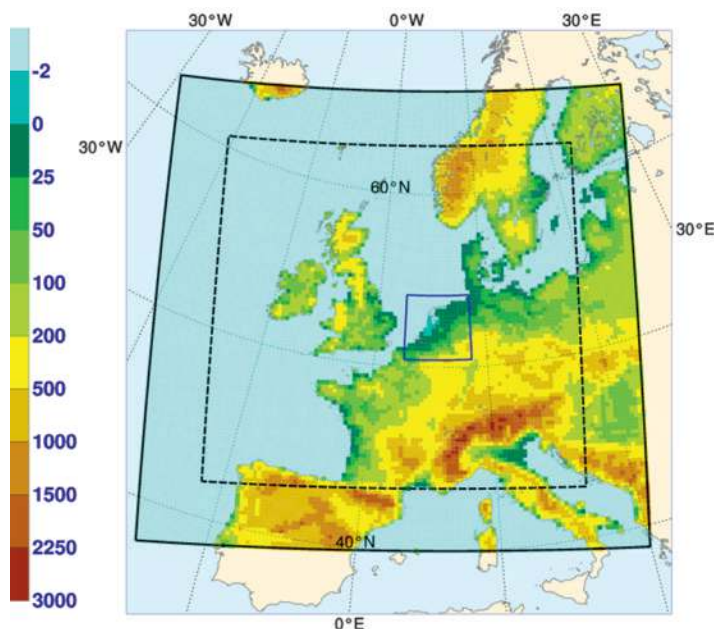
Fɪɢ. 6. **Domain nesting. The solid black line represents the domain of RACMO, spanning 126 × 120 points in the horizontal directions; the dashed line indicates the transition between inner domain and border region. The solid blue line depicts the domain of GALES. Colors indicate the surface level in meters.**

APPENDIX: SETUP. The initial and lateral boundary conditions were provided by the Regional Atmospheric Climate Model (RACMO) that was developed at the KNMI (Royal Netherlands Meteorological Institute). RACMO is a hydrostatic model utilizing the physics parameterizations from the ECMWF model, extensively described in van Meijgaard et al. (2008). The domain of RACMO is illustrated in Fig. 6. The LES domain is approximately from 2 ... 54.09° N (400 km x ...

The original LES ... Heus et al. (2010); ... here. Following Böin... proximation was replaced ... to account for density variations in the vertical direction, and a simple, single-moment microphysics scheme is employed. These modifications ... of the domain to the 15-km ...

The vertical resolution increases ... from $\Delta z = 40$ m at the surface to 70 ... the domain ... of $\Delta x = \Delta y$ ... ary layer ... characteriz... which corr... $Re_l = (z_i / \Delta$ ... first- and s... ... Patton 2011), but higher effective Reynolds ... are desirable.

During the LES runs, we retained periodic boundary conditions to allow developing fluctuations to reenter on the opposite ... border ... around the domain, the ... fixed ... RACMO. This procedure, outlined as follows, ... provides boundary layer conditions to the ... while retaining turbulent fluctuations. ... variables $\{u,v,\theta_l,q_t\}$: longitudinal velocity, lati-... velocity, liquid water potential temperature, and specific humidity, respectively, are relaxed toward the RACMO state in the border regions. First, the average of a variable $\phi \in \{u,v,\theta_l,q_t\}$ is calculated within ... region:

$$\langle\phi\rangle_b = \frac{\iint_S f\phi\,dxdy}{\iint_S f\,dxdy}$$

where S is a single subdomain of 1 GPU (25 km x 25 km), and f is a factor that creates a smooth transition between border region and inner domain: $f = 1$ at the boundary and $f = 0$ well into the inner domain. The tendency due to the boundary conditions (bc) is given by

$$\left[\frac{\partial\phi}{\partial t}\right]_{bc} = -\frac{f}{\tau}\left(\langle\phi\rangle_b - \phi^R\right)$$

where the superscript R denotes RACMO's value of the variable. This formulation preserves turbulence while allowing for heterogeneous boundaries. The border region comprises a 12-km-wide edge around the square domain, and the relaxation time scale τ is 60 s, roughly the time scale at which high winds move through the border region. Another 12 km is removed from the simulation results to allow turbulence to adapt to the new mean state. Alternatively, one might provide in/outflow boundary conditions, but also here the supply of turbulent fluctuations remains a complex problem (e.g., Mirocha et al. 2013).

The conservation equation for an LES filtered variable $\tilde\phi$ becomes:

$$\frac{\partial\tilde\phi}{\partial t} = -\underbrace{\frac{1}{\rho_0}\nabla\cdot\rho_0\tilde{\mathbf{u}}\tilde\phi}_{\text{resolved}} - \underbrace{\frac{1}{\rho_0}\nabla\cdot\rho_0\widetilde{\mathbf{u}^s\phi^s}}_{\text{sub grid}}$$
$$-\underbrace{w^R\frac{\partial\tilde\phi}{\partial z}}_{\text{subsidence}} - \underbrace{\frac{f}{\tau}\left(\langle\phi\rangle_b - \phi^R\right)}_{\text{bc}} + \underbrace{S_\phi}_{\text{sources}} \qquad (A1)$$

where $\mathbf{u} = (u,v,w)$ is the velocity vector.

The first two terms on the right-hand side represent the advection of resolved and subfilter-scale turbulence, respectively. Subfilter-scale motions in GALES, denoted with the superscript s, are treated through eddy viscosity/diffusivity fluxes, modeled as a function of the subfilter-scale turbulent kinetic energy; $\rho_0 = \rho_0(z)$ is a base density profile that is dependent on height only.

The third and fourth term are meant to represent the effects of large-scale vertical and horizontal transport, respectively. The third term models subsidence, which acts as a slow downward (or upward) force. By using RACMO's state for the lateral boundary conditions (fourth term), the LES is effectively provided with large-scale horizontal advection into the LES domain.

For $\phi \in \{\theta_l,q_t\}$, the source term S_ϕ includes effects from precipitation and radiation. Since no GPU-accelerated radiation module was available to GALES at the time of the simulation, we employed RACMO's radiative tendencies for the latter.

The momentum equations are essentially equal to Eq. (A1) if S_ϕ is taken to include tendencies due to pressure fluctuations (Heus et al. 2010) and the ageostrophic acceleration $2\Omega \times (\tilde{\mathbf{u}} - \mathbf{u}_g^R)$, with $\mathbf{u}_g^R$ the geostrophic wind from RACMO.

Last, GALES's surface model was modified to handle a heterogeneous roughness map at 100-m resolution, as well as a high-resolution land–sea map (converted from Farr et al. 2007). Based on the latter, the surface model switches between a prescribed sea surface temperature formulation and a land surface model that accepts the net radiative surface flux as input.

FOR FURTHER READING

Arakawa, A., J.-H. Jung, and C.-M. Wu, 2011: Toward unification of the multiscale modeling of the atmosphere. *Atmos. Chem. Phys.,* **11,** 3731–3742, doi:10.5194/acp-11-3731-2011.

Betts, A., and C. Jakob, 2002: Study of diurnal cycle of convective precipitation over Amazonia using a single column model. *J. Geophys. Res.,* **107,** 4732, doi:10.1029/2002JD002264.

Böing, S. J., H. J. J. Jonker, A. P. Siebesma, and W. W. Grabowski, 2012: Influence of the subcloud layer on the development of a deep convective ensemble. *J. Atmos. Sci.,* **69,** 2682–2698, doi:10.1175/JAS-D-11-0317.1.

Bougeault, P., 1981: Modeling the trade-wind cumulus boundary layer. Part I: Testing the ensemble cloud relations against numerical data. *J. Atmos. Sci.,* **38,** 2414–2428, doi:10.1175/1520-0469(1981)038<2414:MTTWCB>2.0.CO;2.

Bou-Zeid, E., C. Meneveau, and M. Parlange, 2005: A scale-dependent Lagrangian dynamic model for large eddy simulation of complex turbulent flows. *Phys. Fluids,* **17,** 025105, doi:10.1063/1.1839152.

Chlond, A., 1992: Three-dimensional simulation of cloud street development during a cold air outbreak. *Bound.-Layer Meteor.,* **58,** 161–200, doi:10.1007/BF00120757.

Deardorff, J. W., 1972: Numerical investigation of neutral and unstable planetary boundary layers. *J. Atmos. Sci.,* **29,** 91–115, doi:10.1175/1520-0469(1972)029<0091:NIONAU>2.0.CO;2.

de Roode, S. R., P. G. Duynkerke, and H. J. J. Jonker, 2004: Large-Eddy Simulation: How large is large enough? *J. Atmos. Sci.,* **61,** 403–421, doi:10.1175/1520-0469(2004)061<0403:LSHLIL>2.0.CO;2.

Dufresne, J. L., and S. Bony, 2008: An assessment of the primary sources of spread of global warming estimates from coupled atmosphere-ocean models. *J. Climate,* **21,** 5135–5144, doi:10.1175/2008JCLI2239.1.

European Centre for Medium-Range Weather Forecasts, 2014: Evolution of the ECMWF analysis and forecasting system. [Available online at www.ecmwf .int/products/data/operational_system/evolution.]

Farr, T. G., and Coauthors, 2007: The Shuttle Radar Topography Mission. *Rev. Geophys.,* **45,** RG2004, doi:10.1029/2005RG000183.

Hatlee, S., and J. Wyngaard, 2007: Improved subfilter-scale models from the HATS field data. *J. Atmos. Sci.,* **64,** 1694–1705.

Heus, T., and Coauthors, 2010: Formulation of the Dutch Atmospheric Large-Eddy Simulation (DALES) and overview of its applications. *Geosci. Model Dev.,* **3,** 415–444, doi:10.5194/gmd-3-415-2010.

Jonker, H. J. J., M. van Reeuwijk, P. P. Sullivan, and E. G. Patton, 2013: On the scaling of shear-driven entrainment: A DNS study. *J. Fluid Mech.,* **732,** 150–165, doi:10.1017/jfm.2013.394.

Khairoutdinov, M., and D. Randall, 2006: High-resolution simulation of shallow-to-deep convection transition over land. *J. Atmos. Sci.,* **63,** 3421–3436, doi:10.1175/JAS3810.1.

——, S. K. Krueger, C.-H. Moeng, P. A. Bogenschutz, and D. A. Randall, 2009: Large-Eddy Simulation of maritime deep tropical convection. *J. Adv. Model. Earth Syst.,* **1,** 15, doi:10.3894/JAMES.2009.1.15.

Lilly, D. K., 1962: On the numerical simulation of buoyant convection. *Tellus,* **14 (2),** 148–172, doi:10.1111/j.2153-3490.1962.tb00128.x.

Mirocha, J., G. Kirkil, E. Bou-Zeid, F. Chow, and B. Kosović, 2013: Transition and equilibration of neutral atmospheric boundary layer flow in one-way nested large-eddy simulations using the Weather Research and Forecasting Model. *Mon. Wea. Rev.,* **141,** 918–940, doi:10.1175/MWR-D-11-00263.1.

Miura, H., M. Satoh, T. Nasuno, A. T. Noda, and K. Oouchi, 2007: A Madden-Julian oscillation event realistically simulated by a global cloud-resolving model. *Science,* **318,** 1763–1765, doi:10.1126/science.1148443.

Miyamoto, Y., Y. Kajikawa, R. Yoshida, T. Yamaura, H. Yashiro, and H. Tomita, 2013: Deep moist atmospheric convection in a subkilometer global simulation. *Geophys. Res. Lett.,* **40,** 4922–4926.

Moeng, C.-H., J. Dudhia, J. Klemp, and P. Sullivan, 2007: Examining two-way grid nesting for large eddy simulation of the PBL using the WRF Model. *Mon. Wea. Rev.,* **135,** 2295–2311, doi:10.1175/MWR3406.1.

Molinari, J., and M. Dudek, 1992: Parameterization of convective precipitation in mesoscale numerical models: A critical review. *Mon. Wea. Rev.,* **120,** 326–344, doi:10.1175/1520-0493(1992)120<0326:POCPIM>2.0.CO;2.

Perot, J., and J. Gadebusch, 2009: A stress transport equation model for simulating turbulence at any mesh resolution. *Theor. Comput. Fluid Dyn.,* **23,** 271–286, doi:10.1007/s00162-009-0113-x.

Rauser, F., 2014: High definition clouds and precipitation for advancing climate prediction. [Available online at http://hdcp2.zmaw.de.]

Rossow, W. B., G. Tselioudis, A. Polak, and C. Jakob, 2005: Tropical climate described as a distribution of weather states indicated by distinct mesoscale cloud property mixtures. *Geophys. Res. Lett.,* **32,** L21 812, doi:10.1029/2005GL024584.

Schalkwijk, J., E. Griffith, F. H. Post, and H. J. J. Jonker, 2012: High-performance simulations of turbulent clouds on a desktop PC: Exploiting the GPU. *Bull. Amer. Meteor. Soc.,* **93,** 307–314, doi:10.1175/BAMS-D-11-00059.1.

Shuman, F. G., 1989: History of Numerical Weather Prediction at the National Meteorological Center. *Wea. Forecasting,* **4,** 286–296, doi:10.1175/1520-0434(1989)004<0286:HONWPA>2.0.CO;2.

Simmons, A. J., D. M. Burridge, M. Jarraud, C. Girard, and W. Wergen, 1989: The ECMWF medium-range prediction models development of the numerical formulations and the impact of increased resolution. *Meteor. Atmos. Phys.,* **40,** 28–60, doi:10.1007/BF01027467.

Skamarock, W., and Coauthors, 2008: A description of the advanced research WRF Version 3. Tech. Rep. Tech. Note NCAR/TN-475+STR, NCAR, 113 pp.

——, J. Klemp, M. Duda, L. Fowler, S.-H. Park, and T. Ringler, 2012: A multiscale nonhydrostatic atmospheric model using centroidal voronoi tesselations and C-grid staggering. *Mon. Wea. Rev.,* **140,** 3090–3105, doi:10.1175/MWR-D-11-00215.1.

Slingo, J., and T. Palmer, 2011: Uncertainty in weather and climate prediction. *Philos. Trans. R. Soc., A,* **369,** 4751–4767, doi:10.1098/rsta.2011.0161.

Smagorinsky, J., 1974: Global atmospheric modeling and the numerical simulation of climate. *Weather and Climate Modification,* John Wiley & Sons, 633–686.

Sommeria, G., 1976: Three-dimensional simulation of turbulent processes in an undisturbed trade wind boundary layer. *J. Atmos. Sci.,* **33,** 216–241, doi:10.1175/1520-0469(1976)033<0216:TDSOTP>2.0.CO;2.

Stevens, B., and S. Bony, 2013: Climate change: What are climate models missing? *Science,* **340,** 1053–1054, doi:10.1126/science.1237554.

Strohmaier, E., J. J. Dongarra, H. W. Meuer, and H. D. Simon, 2005: Recent trends in the marketplace of high performance computing. *Parallel Comput.,* **31,** 261–273, doi:10.1016/j.parco.2005.02.001.

Sullivan, P., and E. Patton, 2011: The effect of mesh resolution on convective boundary layer statistics and structures generated by large-eddy simulation. *J. Atmos. Sci.,* **68,** 2395–2415, doi:10.1175/JAS-D-10-05010.1.

van Meijgaard, E., L. H. van Ulft, W. J. van de Berg, F. C. Bosveld, B. J. J. M. van den Hurk, G. Lenderink, and A. P. Siebesma, 2008: The KNMI regional atmospheric climate model RACMO version 2.1. Tech. Rep. 320, KNMI, 43 pp.

Wyngaard, J. C., 2004: Toward numerical modeling in the "terra incognita." *J. Atmos. Sci.,* **61,** 1816–1826., doi:10.1175/1520-0469(2004)061<1816:TNMITT>2.0.CO;2.

Radar and Atmospheric Science:
A Collection of Essays in Honor of David Atlas

Edited by Roger M. Wakimoto and Ramesh Srivastava

This monograph pays tribute to one of the leading scientists in meteorology, Dr. David Atlas. In addition to profiling the life and work of the acknowledged "Father of Radar Meteorology," this collection highlights many of the unique contributions he made to the understanding of the forcing and organization of convective systems, observation and modeling of atmospheric turbulence and waves, and cloud microphysical properties, among many other topics. It is hoped that this text will inspire the next generation of radar meteorologists, provide an excellent resource for scientists and educators, and serve as a historical record of the gathering of scholarly contributions honoring one of the most important meteorologists of our time.

Radar and Atmospheric Science:
A Collection of Essays in Honor of David Atlas

Aug 2003. Meteorological Monograph Series, Vol. 30, No. 52; 270 pp, hardbound; ISBN 1-878220-57-8; AMS code MM52.

Price $75.00 list/$55.00 member

To place an order point your Web browser to www.ametsoc.org/amsbookstore

AMS BOOKS

RESEARCH ◆ APPLICATIONS ◆ HISTORY

RADAR AND ATMOSPHERIC SCIENCE:
A COLLECTION OF ESSAYS IN HONOR OF
DAVID ATLAS

Edited by
Roger M. Wakimoto
Ramesh C. Srivastava

Published by the American Meteorological Society