

Interactive comment on “Towards European-Scale Convection-Resolving Climate Simulations” by David Leutwyler et al.

David Leutwyler et al.

david.leutwyler@env.ethz.ch

Received and published: 30 August 2016

Thank you for your positive and encouraging feedback. The version of COSMO, capable of exploiting GPUs, has been - and still is - a big effort by a large team of contributors and we appreciate that you welcome our work.

This is a clearly written manuscript describing world leading efforts to apply high-performance computing to climate studies. With some sharpening of focus, it could make a worthwhile contribution to the literature. Specifically I recommend expanding and deepening section 5, and shortening and making more specific sections 3 and 4.

Similar issues were also raised by Reviewer 3. We revised and shortened the text of sections 3 and 4 and also added more section titles. To shorten the descriptive parts of the manuscript in section 3 and 4, we removed Figure 4 and we moved one of the cases

C1

displayed in former Figure 6 to the supplementary material. Furthermore we removed the panels, displaying relative humidity, in former Figure 11. The corresponding text has also been shortened significantly (most notably Sections 3.1.1 Kyrill Evolution, 3.1.2 Representation of Precipitation along Cold Fronts and 3.1.3 Representation of a Meso-scale Low).

In addition, we substantially expanded section 5:

(1) We clarified the “socket metric” and the implications for the presented experiment (including a new figure, Figure 11 in the revised version). The respective paragraphs now read:

“The full strong-scaling experiment corresponds to a 24 h simulation on a domain of $1536 \times 1536 \times 60$ grid points. Input for this simulation consists of the lateral boundary conditions at hourly resolution, amounting to about 120 GB for the whole simulation. Additionally an output workload consisting of about 6 GB is written to the file system. All performance results have been obtained on a heterogeneous Cray XC30 system, located at the Swiss National Supercomputing Centre (CSCS) in Lugano, Switzerland (Piz Daint). The Piz Daint supercomputer consists of a heterogeneous node architecture with an eight-core Intel Xeon E5-2670 CPU and an NVIDIA Tesla K20X GPU per node (Figure 11), and Cray’s Aries interconnect using a three-level dragonfly topology to connect the compute nodes. To normalize the performance metrics, they are defined as per socket. In the case of our configuration (Piz Daint, Figure 11), a socket corresponds to either an eight-core Xeon CPU or an NVIDIA K20x GPU.

A socket is the electrical component that provides the connection between the circuit board and the chip sitting on top of it. The advantage of the per-socket metric is its flexibility across architectures, which also allows comparing with individual sockets on a multi-GPU node (fat node). On a fat node, a socket still hosts only a single GPU chip, even if multiple GPU sockets are installed on a PCI express card or on a node. However, for the node configuration found in Piz Daint, this metric is a bit unfair towards

C2

the multi-core systems, since GPUs (today) still need an accompanying CPU hosting the operating system and instructing the GPU. With the socket-based metric, we do not account for that additional CPU. Another metric would be node-to-node comparison, assuming that a node can either consist of one CPU and a GPU, or two CPUs. For such a configuration, the second option would be fairer for the multi-core architecture. In general, node-to-node comparison is useful to compare the various possible node configurations one may find in a supercomputer. However, we believe that for the current study the per-socket performance metric is more useful than node-to-node comparisons, also because nowadays fat-nodes are commercially available."

(2) We elaborate differences in weak-scaling experiments w. r. t. global models by adding the following paragraph: "The weak-scaling approach used here is slightly different for weak-scaling experiments with global simulations, because in these experiments the domain size can not be varied, except by shrinking and expanding the size of the planet. In some global experiments the grid-spacing is varied while keeping the time step constant (Wehner et al., 2011)."

(3) We discuss potential differences in the weak-scaling behavior w. r. t. the employed numerical schemes in Section 5.3 (Assessment). We added the following paragraph in the end of the section: "Whether the above assessment can be transferred to GCMs, also depends upon the time-stepping algorithm and its implementation. Many global non-hydrostatic models invoke semi-Lagrangian or spectral approaches (where the total communication costs increases faster than the number of gridpoints). In such cases, perfect weak scaling will be more difficult to achieve than with the current split-explicit scheme."

For further details on these changes, please refer to the revised/marked-up manuscript for the details.

I suggested minor revisions as my comments and critique should be sufficiently clear for the editor to adjudicate on his/her own, not as an indication of their importance. But

C3

also for this reason I would not be inclined to re-review a revised manuscript. This is something the editor can do.

Specific Comments:

1. *The conclusion (line 585) makes a concise point of the additional mesoscale structure that becomes representable as the grid is refined to ca 2 km. The manuscript itself is too descriptive and qualitative, and the added value of the 2 km simulations should be more concisely and clearly presented.*

Based on your suggestion and the suggestions of Reviewer 3, we added more quantitative information in section 3 (Kyrill case). Besides sharpening the text in general, we added the following:

(1) We provide an estimate on the width of the cold frontal rain bands:

"In CTRL12, the front is split into successive precipitation bands with maximum precipitation rates up to 20 mm/h. CTRL2 additionally features small-scale embedded convection located in the vicinity and along the front, and a more coherent organization. The frontal rainbands (precipitation > 5 mm/h) are typically 30-40 km wide in CTRL12, and substantially narrower in CTRL2 (8-10 km)."

(2) After moving one of the cases in former Figure 6 to the supplementary material, we added another row of panels for the remaining case (now Fig. 5). The refined display allows us to establish the representation of narrow cold-frontal rainbands. The corresponding text reads:

"The distinct narrow cold-frontal rainbands seen in the bottom-right panel of Figure 5 are of distinctly convective origin. They are associated with precipitation rates >20 mm/h, located on the leading edge of the fronts, and aligned with the cold front in an oblique angle. These systems have been extensively discussed in the literature (Houze, 2014), and studied using (airborne) radar (e. g. Jorgensen et al. 2003). We expect differences in location and intensity, due to the ability of CTRL2 to explicitly

C4

resolve the underlying dynamical processes.”

Short simulations with atmospheric models in climate mode are of limited predictability and detailed verification of the discussed cases would therefore be dubious. Furthermore, the observational reference of the cases discussed is too weak to allow for detailed verification of the underlying processes.

2. The most interesting aspect of the manuscript was the proof of concept and computational perspective of the outlook (section 5). Emphasizing the scaling performance, and where and how this might be expected to change either based on different implementations or hardware changes could, when combined with point 1, more sharply guide the community.

We are very glad that you consider section 5 useful for the community. We substantially expanded it, including a new paragraph on the weak-scaling behavior of global models and added a new figure (current Figure 11). As outlined in more detail above, we clarified the “socket metric” and the implications for the presented experiment. Furthermore we elaborate differences in weak-scaling experiments w. r. t. global models and w. r. t. numerical schemes.

However, the question how the (stencil) implementations scale on other hardware architectures, than those two considered in the current study, is a rather complex topic and would require a dedicated manuscript.

Details:

line 8: COSMO (spell out)

Changed according to the GMD guidelines.

The sentence now reads: “One of the first atmospheric models that has been fully ported to these architectures is the Consortium for Small-scale Modeling model (COSMO).”

C5

line 21: What makes COSMO a climate model? I don't think this is such and appropriate use of the phrase.

We refer to our version of COSMO as a “regional climate model” for two reasons: First, the model includes representations of the couplings with the land-surface, in particular with soil moisture and snow cover. Second, the model enables simulations over climatological time periods (at least a decade), and can thus be validated not only in an NWP context, but also with respect to climatological distributions. On a more technical level, for climate-scale simulations with COSMO, several enhancements not needed in NWP mode, such as a time-dependent leaf-area index or time-dependent sea surface temperatures, are needed. Many of these additional climate features are maintained by the COSMO-CLM community. However, since the COSMO v4.19 is not an official CLM version, we refrain from calling our model version COSMO-CLM for now.

The sentence has been changed to: “With the COSMO model, we now use a weather and climate model that has all the necessary modules required for real-case convection-resolving regional climate simulations on GPUs.”

line 29: Fine, but I also think this understates what might be an important role for the mesoscale.

We also think that convection-resolving models provide substantial insight about meso-scale processes. It should therefore not miss out here.

We modified the first paragraph accordingly. It now reads: “The inadequate representation of clouds and moist convection represents a major challenge of state-of-the-art climate models (Stevens and Bony, 2013). An important component of the problem are the scale interactions between small-scale turbulent and convective processes at scales around and below 1 km, and larger-scale/meso-scale weather systems at scales around O(10 km-1000 km). Within these scale interactions, individual convective cells may organize into meso-scale weather systems such as squall lines or meso-scale

C6

convective systems. Current global and regional climate models typically operate at grid spacings on the order of 10-300 km, and are thus unable to explicitly represent many of these interactions."

line 32: I never think of Dai and Trenberth as a reference for the uncertainties, or approximations required in cumulus parameterization. line 40: A case could be made for mentioning super-parameterization here

We would like to address these two issues together by citing Randall et al. (2003) instead. They provide a nice overview of cloud-related uncertainties, and propose the super-parametrization concept.

line 50: 'partly resolved' is 'partly' true. What about the stable boundary layer, or clouds that are only a few hundreds of meters, or do you mean this to apply to deep convective systems. More clarity would be helpful.

Yes, some clarifications could help indeed. We changed partly resolved to underresolved and added further explanations.

The first part of the paragraph now reads: "While the convection-resolving approach shows very promising results, turbulent and convective motions are still underresolved (Wyngaard, 2004). Grid spacings of O(1 km) are comparable to the size of the particularly energetic convective eddies in the planetary boundary layer (Zhou et al., 2014). At this resolution, shallow clouds still need to be parametrized and deep-convective clouds tend to be too large, too laminar, too vicious and too widely spaced apart (Clark et al., 2016). Using numerical simulations of an idealized squall line, Bryan et al. (2003) showed ..."

line 60: "narrow"? Perhaps "fine" would be a better choice.

Yes, it is. Changed accordingly.

line 146: How adequate is the Tiedtke references? Is this what COSMO really uses . . . I find it hard to believe that the implementation is unchanged from what Tiedtke

C7

describes for implementation at ECMWF.

We think it is still adequate to quote the Tiedtke (1989) reference, although the version used in COSMO has indeed been adapted from the original implementation. Some of these changes are described in the COSMO documentation, some are fairly recent and have not yet been documented thoroughly.

Line 152: Is it necessary to introduce the DESL acronym? Line 168: Same goes for GCL

We removed these acronyms from the manuscript, since they are indeed not used later on.

Line 228: I am not sure I would lead by saying how you don't do things.

Based on the suggestion, we substantially altered the structure of the respective subsection. Please see the revised version of the manuscript.

*Line 240: Traditionally abbreviations, such as "srf" in mathematical text are typed in the roman font. Variables take on the italic font. This helps distinguish between "srf" meaning surface, as opposed to 'srf' meaning the entropy times the radius times the free energy. . . or whatever s*r*f might be.*

We agree and have consulted the SI brochure 8 (mentioned on the GMD website). We changed the symbols accordingly.

Line 255: I appreciate your humanity; but describing this as 'sad' might assign unintended value to the situation. For instance, some might argue that 'sad' is too passive a voice. I would simply state that it is estimated but there were X storm related casualties and damage estimates were Y (ref).

We decided to remove this sentence from the manuscript. We chose the Kyrill case because it has been well assessed by Fink et al. (2009) and Ludwig et al. (2015). Actually, we do not assess the impact of the storm itself.

C8

Line 272: Check the style guide but I am not used to seeing square brackets to delineate ellipsis.

We found no recommendation in the GMD guides, but it seems common practice to use round brackets. We changed it here and in the rest of the manuscript accordingly.

Line 286: Is "figure 3" a proper name, i.e., "Figure 3"

Yes. We changed all the occurrences of that term accordingly.

Line 315: Is there no way to look directly at station data so as to infer the core pressures? I guess this depends on where the core is.

While composing this manuscript we tried to find station data. For instance, we also tried to infer station data from the Synop- and Bodenwetterkarten of DWD. At least in our sources, there are not sufficient stations available that are located close enough to the storm core, even after it made landfall.

Figure 7: Why not continue to show CTRL 50

Good idea. We have added panels with results from CTRL50 and adapted the text accordingly (this is now in Fig.6, as we have dropped the previous Fig.4). At the same time we also reduced visual clutter in the figure, stemming from the coordinate labels. At this occasion we also corrected a small bug in the display: In the initial version, smoothing was applied to the geopotential height contours in the panels of precipitation and temperature, which was removed. This bugfix makes the geopot. contours slightly noisy, but makes all three rows of the display consistent, and does not affect the conclusions drawn from the analysis.

Line 346: Is the 'polar low' designation correct? Would be good to be specific and connect better to the literature?

When composing the initial version of this manuscript we used the following definition: "A polar low is a small, but fairly intense maritime cyclone that forms poleward of the

C9

main baroclinic zone (the polar front or other major baroclinic zone). The horizontal scale of the polar low is approximately between 200 and 1000 kilometres and surface winds near or above gale force."

Erik A. Rasmussen and John Turner (2003). *Polar Lows*. doi: 10.1017/CBO9780511524974

Actually we do not assess the relevant forcing mechanisms in detail and for this study it is not particularly relevant if it this particular eddy can be classified as polar low or not. We therefore replaced the term "polar low" by "meso-scale vortex", which is a less restrictive decision.

Line 401: But isn't one advantage of the 2km model that it captures orographic influences, and in this case comparing the valley totals to the mean would help quantify this statement?

It has indeed been shown that the added detail, provided by the increased resolution, improves the representation of orographic influence (see Ban et al. 2015, and Prein et al., 2015 for a review). As explained in the text and the references therein, the low station density and uneven distribution introduces substantial uncertainties in areas with high spatial variability. The sampling uncertainty will manifest itself as long as (low resolution) gridded data is used. Direct validation with station data could help circumvent this problem

The goal of Figure 8 (revised version) is to validate the continental scale pattern. Thoroughly validating valley and mountain-peak totals would involve using station data and high-resolution data sets. A study with a more local-scale focus would be better suited for such a validation. Moreover, as you suggested in the introductory comment, the particular section in the manuscript should be shortened rather than extended.

Line 405: Isn't the ability of convective resolving/permitting simulations to better simulate the diurnal cycle by now an old result?

C10

Yes. That is why we state the following a few lines below:

L410: "It has previously been shown for smaller domains (Kendon et al., 2012; Ban et al., 2014) that the behavior of the convection-resolving model fits observation much better. Our results are qualitatively consistent with these studies, although the differences in daily precipitation statistics are larger for our simulation."

A shortcoming of the simulation used in the address section of the manuscript is that it is only a few months long. It would be therefore be questionable to conduct detailed validation on local scales using observational precipitation data sets. Therefore, we mostly perform model-to-model comparisons. Our concept is to use well established results, from simulations on smaller domains (e. g. Hohenegger et al. 2009 or Ban et al., 2014), to show that their findings also apply to simulations on European-scale domains.

Line 430: Maybe an earlier Tompkins and Craig reference is warranted in the discussion of cold pools?

It definitely is. We think that Tompkins (2001) should be suited best.

Line 460: One should distinguish between weak scaling at fixed resolution and weak scaling more generally. The distinction is that the temporal dimension does not exhibit weak scaling.

That could indeed be useful information for the reader, especially since weak-scaling experiments in regional and global models are typically performed slightly differently. With a global model, an experiment similar to the one performed here, would be shrinking and expanding the size of the planet while proportionally adapting the number of nodes used. However, weak-scaling experiments are rarely performed this way. Instead, typically the resolution is decreased, while keeping the time step constant. When performing a weak-scaling experiment with a regional model, the grid spacing, time step and the time to solution can remain the same across all domain sizes. For

C11

strong scaling, on the other hand, only the number of grid points per node changes, while the time dimension cannot be scaled.

To clarify the different approaches, we adapted the first part of the paragraph the following way:

"What are the computational requirements to perform a convection-resolving simulation on the European scale? Here we restrict the analysis to two key performance metrics:

1. Strong scaling: The achievable time to solution for a fixed simulation domain, fixed grid spacing and domain size, while increasing the computational resources. For linear scaling, the time to solution will increase inverse proportionally to the computational resources, allocated to the problem. Here the time step (which is constrained by the grid spacing through the Courant stability criterion) can be kept constant and hence the computational task has a fixed size
2. Weak scaling: The achievable time to solution when the domain size is increased proportionally with the computational resources, while keeping the grid spacing and the time step fixed. For linear scaling, the time to solution would remain the same for all domain sizes.

The weak-scaling approach used here is slightly different for weak-scaling experiments with global simulations, because in these experiments the domain size can not be varied, except by shrinking and expanding the size of the planet. In some global experiments the grid-spacing is varied while keeping the time step constant at the value required by the simulation with the finest grid (e. g. Wehner et al., 2011)."

Line 490: This discussion would benefit from some reflection on what the trade offs are of the different ways of comparing things.

This point was also raised by other reviewers. We added further explanations and it now reads:

C12

"(...) To normalize the performance metrics, they are defined as per socket. In the case of our configuration (Piz Daint, Figure 11), a socket corresponds to either an eight-core Xeon CPU or an NVIDIA K20x GPU.

A socket is the electrical component that provides the connection between the circuit board and the chip sitting on top of it. The advantage of the per-socket metric is its flexibility across architectures, which also allows comparing with individual sockets on a multi-GPU node (fat node). On a fat node, a socket still hosts only a single GPU chip, even if multiple GPU sockets are installed on a PCI express card or on a node. However, for the node configuration found in Piz Daint, this metric is a bit unfair towards the multi-core systems, since GPUs (today) still need an accompanying CPU hosting the operating system and instructing the GPU. With the socket-based metric, we do not account for that additional CPU. Another metric would be node-to-node comparison, assuming that a node can either consist of one CPU and a GPU, or two CPUs. For such a configuration, the second option would be fairer for the multi-core architecture. In general, node-to-node comparison is useful to compare the various possible node configurations one may find in a supercomputer. However, we believe that for the current study the per-socket performance metric is more useful than node-to-node comparisons, also because nowadays fat-nodes are commercially available."

Line 500: Why compare 64x64 with 2620, as it forces the reader to do the multiplication. I prefer to call 64x64 4096.

Indeed. We changed that number according to your suggestion.

Line 510: Some additional performance metrics, like time per price, or time per power would be informative, even if only qualitative.

The issues raised here are among the key questions when pushing towards global kilometer-scale models. While a qualitative result would indeed be very useful to scientists in the atmospheric domain, others would probably not appreciate us estimating these numbers in an ad-hoc way. Energy to solution experiments have to be carried

C13

out with great care and would require designing a dedicated experiment and study.

A video from an online presentation at SC13 from Thomas Schulthess et al., which includes an analysis of energy to solution can be found here: <https://www.youtube.com/watch?v=X5PqyfXc9pAt=21m20s>

Line 530: It may be mention that many CORDEX simulations are performed with very old models (REMO) whose parallel efficiency is quite small.

We hesitate to criticize other models and modeling groups without performing a thorough performance assessment ourselves.

Line 556: I would prefer to phrase this as: "what does this mean for global simulations"

We changed the term accordingly.

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-119, 2016.

C14