

Contents

- Welcome to our new RSE!
- New component structure
- natESM on podcast
- New newsletter category
- deep dive with Clara Bayley
- Sprint status

New component structure

We are pleased to share the new component structure within the natESM system with you:

- 1. Core components:** Fundamental models for basic Earth system modeling (atmosphere, ocean, land).
- 2. Extended core components:** Comprehensive models for Earth-system interactions (e.g., atmospheric chemistry, ocean biogeochemistry, land ice, socioeconomics).
- 3. Infrastructure Components:** Tools and interfaces that connect core and extended core components, ensuring seamless functionality.
- 4. Impact Components:** Models for impact modeling, detailing interactions and consequences within the system.
- 5. Optional Components:** Community-suggested models for future integration, visible for exploration but not part of the main system. These can become extended core components if they meet our acceptance criteria as outlined in the [natESM strategy](#).

You're welcome to send your feedback as we navigate this process together and find our way forward. Visit our [system on GitLab](#), where we just added numerous optional components with detailed information on contacts, code languages, links, literature, projects etc.. If you'd like to advertise your own, please fill out our [Google Form](#).

New team member at DKRZ



We are very happy to welcome our new RSE, [Aparna Devulapalli](#)! Aparna has a background in computational science and engineering with dual master's degrees from esteemed universities in Germany and Switzerland. She is passionate about using her skills to develop advanced scientific technologies in HPC, focusing on atmospheric sciences and climate research. Aparna has extensive experience working with HPC machines as part of her academic courses and projects. We are excited to have her on board!

natESM on podcast

The podcast "[Code for Thought](#)," which covers software, engineering, research, and more, has released an episode about the natESM project. In this episode, Peter Schmidt talks with Iris Ehlert and Hendryk Bockelmann about the natESM project, our strategy, and other insights. You can [listen to the podcast episode here](#) (German only!).

New newsletter category: deep dive

We are excited to introduce a new category in our newsletter: deep dive - personal stories from natESM trailblazers. This section will occasionally provide in-depth insights into the inspiring journeys of natESM community members. For our first deep-dive story, we interviewed [Clara Bayley from MPI-M](#) to discuss her experiences with the CLEO sprint (next double page).

Register for November training

You're welcome to register for our training taking place on November 5-6 in Jülich. The training will cover performance analysis and GPU programming (CUDA, OpenACC, OpenMP, Kokkos). You can find the [preliminary agenda here](#). For more details and to register, visit our [event page](#).

deep dive

personal stories from natESM trailblazers

For our first deep-dive story, PhD candidate [Clara Bayley](#) from MPI-M reveals how the CLEO sprint transformed the course of her PhD journey. We connected via Zoom just after she arrived in Korea for a workshop and conference. Despite the exhausting flight, she was happy to share her insights about the sprint. Although she is 7 hours ahead in time, she and the responsible RSE, Wilton Loch from DKRZ, have been communicating through Zoom consistently throughout the sprint, making her temporary absence a non-issue for the ongoing progress of their collaboration.

Interview: Iris Ehlert

natESM: Clara, with about 1.5 months left in your first natESM sprint, could you give us an overview of the goals you set and how far you've come in achieving them?

Clara: Sure! Our sprint has been focused on two main objectives. Firstly, we've been working on coupling my superdroplet model to ICON, which involves exchanging crucial data such as temperature and pressure. We've made significant progress with the one-way coupling—just some final testing left to ensure everything's solid. Secondly, we've been deep into implementing domain decomposition within CLEO. This has included determining the best methods to divide up the workload across different nodes and ensuring smooth data flow for superdroplet motion. Currently, we have a basic model in place that we're refining, with the next big step being setting up parallel output to thoroughly test our domain decomposition.

What were the main challenges you faced during the sprint, and how did working with Wilton help in overcoming them?

To be honest, the sprint has gone incredibly smoothly. We've faced some minor, funny small stuff that we haven't really thought of, like that in the superdroplet model we're using dimensionless variables whereas ICON's fields have dimensions. Realizing that we were passing values about 273 times higher than they should be, led to some funny moments. But these were mostly quick fixes, just a line of code here and there. Our step-by-step approach has been effective in keeping us on track without any major setbacks. The collaboration with Wilton has been incredibly beneficial for me. We've established an excellent communication routine, using Mattermost for quick updates and holding weekly Zoom calls that last anywhere from half an hour to over an hour. This has ensured we're always on the same page and kept everything progressing smoothly. Wilton's work during this sprint has been invaluable. The progress he's made in such a short time would have taken me at least a year due to my limited background knowledge. His expertise and efficiency have provided a tremendous advantage, allowing us to move forward much faster than I could have managed alone.

How important were the kickoff and status meetings in guiding your sprint progress?

Both meetings were really helpful. The kickoff meeting was particularly useful for me to lay out my ideas clearly, which



Clara in her office at the MPI-M. The photo was taken by her mom, who visited her in Hamburg shortly before her departure to Korea.

helped others understand my plans and helped me clarify them for myself. It ensured that everyone had an overview of what was going on, which was beneficial for Wilton as well, as he was interacting with various people throughout the sprint. The status meeting was also valuable. It provided a good opportunity to recap what we had achieved so far, assess how much time we had left, and plan what still needed to be done.

As you near the sprint's end, how are you managing priorities to ensure you meet your goals?

We've established a clear priority list of what we need to accomplish. For example, we've decided to hold off on implementing the two-way coupling for now. Our focus is on getting parallel output working, though we're not aiming for any elaborate many-to-many node output yet. That would be ideal, but it's not our highest priority. Instead, we'll start with getting the output to work on just one node and then focus on implementing the motion aspects. We have a solid plan of what

absolutely needs to be done, what would be nice to have, and what would be perfectionist-level goals. The overarching goals have remained the same; it's just the order in which we've tackled them that has shifted. By the end, the list of what we achieve will likely be scattered among our original goals, but the core objectives remain unchanged. Given what we've done so far and what's left, I doubt we'll complete the two-way coupling, but we're confident about the one-way coupling and the MPI domain decomposition. The one-way (online or offline) coupling is already ready for testing.

Would you apply for another sprint to further develop the two-way coupling?

I'm not sure. It's definitely something I want to do, but I don't know if it would be accepted since it's not a very large project.

It doesn't need to be a large project. You could apply for a shorter sprint, like one for three months.

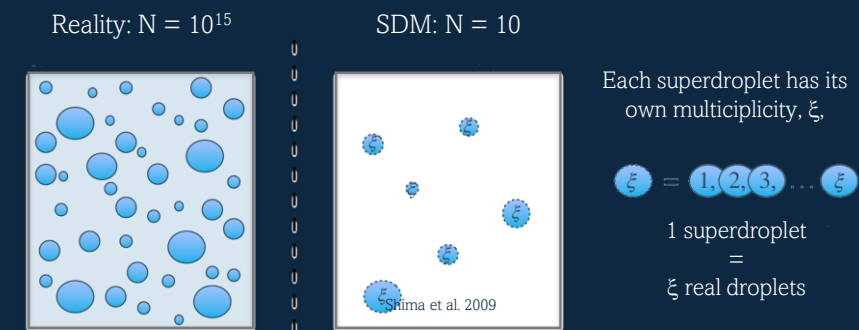
Oh my gosh, I would absolutely do that.

How has the sprint shaped your perspective on collaborative research and technical development?

It's been an all-around positive experience for me. It's been incredibly insightful, especially in understanding how software development is approached by those with a software-development background compared to a scientific one. For instance, in science, we might just share files via Google Drive for convenience, but in my first meeting with Wilton, he asked about naming conventions, and we discussed version control and writing good commit messages. These are things I've heard about but never fully appreciated until now. Learning these technical aspects has been invaluable for me. This sprint was also my first time working directly on code with someone else simultaneously, which was a new and exciting experience.

How do you expect the sprint's outcomes to contribute to natESM community's goals?

I see this sprint as a strong example of how these projects can be effectively structured. Our step-by-step approach demonstrates a workflow that could benefit other sprints. Moreover, in terms of tangible achievements, we're advancing a microphysics scheme for ICON—a widely used model across various institutions like MPI, DWD, and KIT. Once we finalize the two-way coupling, this will provide another robust



A glimpse from the CLEO sprint, adapted from the KickOff presentation.

microphysics option. This development not only enhances capabilities within the community but also showcases modular development for ICON. Looking ahead, our work with MPI domain decomposition illustrates how these techniques can be rapidly implemented. If others in the community face similar resource constraints, our approach could serve as a model for achieving similar goals efficiently.

What benefits have you personally gained from participating in this sprint?

For me, there are two main benefits that stand out. Firstly, the acceleration in progress is significant. When I outline my PhD timeline during panel meetings, I used to allocate six months just to couple to ICON and scale up the model. Now, thanks to the sprint, that timeline has been drastically shortened, freeing up valuable time for more in-depth climate-science research. This has been incredibly empowering. Secondly, the learning experience has been invaluable. I've gained insights into effective collaboration and adopted better coding practices, which I foresee benefiting my research in the long term, including into my postdoc years. These aspects of the sprint are likely to have a lasting impact beyond the immediate time savings in my PhD journey.

What advice would you give to researchers considering applying for a natESM sprint?

Definitely apply. Definitely. Even if you don't have much of a computational background, the RSEs are there to assist you. If you outline what you want to achieve but need help understanding the technical steps, the RSEs guide you on structuring your application with more detail. Having someone kind of help you with saying: Put in more detail like ABC, or like a 1234, these sort of things; I think that massively helps. But the main thing is just to do it. Do a sprint, cause they're so great.

Trainings in July fully booked, check out upcoming events!

The [training on ComIn and YAC](#) (17-18 July) will cover: coupling components to core components, and interactive hands-on sessions with ComIn and YAC. The [training with a spotlight on GPU Programming](#) on LEVANTE Nvidia GPUs (23 July) will cover hands-on sessions on getting started on GPUs and running own kernels. Both trainings are fully booked, but all presentations will be streamed. If you would like to receive the link, please contact Maria Rompe no later than 16 July, noon. The presentations and practical exercises from the hands-on sessions will be made available online after the workshops.

We have listed numerous events on our website, ranging from Summer Schools and Supercomputing Academies to Hackathons. Explore the variety of opportunities available: [natESM Events](#).

Sprint status

SPRINT TITLE	DURATION [MONTHS]	SERVICE DESCRIPTION
ICON-ART Finished	2	Analysis of ART code for GPU porting → Sprint report
ICON-mHM-YAC Finished	6	Online coupling mHM into ICON using YAC → Sprint report
FESOM Finished	6	Port FESOM 2.1 to JUWELS booster and Levante-GPU → Sprint report
ParFlow Finished	6	Port ParFlow to AMD GPUs, Performance Analysis → Sprint report
MESSy Finished	4	Optimize the data transfers between host (CPU) and device (GPU) → Sprint report
ESMValTool Finished	6	Updating remaining non-lazy preprocessor functions to be memory efficient → Sprint report
HAMOCC Finished	6	Concurrent HAMOCC on GPU → Sprint report
MESSy-ComIn Finished	6	Couple MESSy to ICON via the ICON Community Interface → Sprint report
LAGOOn Reporting	6	Develop concept and provide first implementation of Lagrangian-transport-modeling framework
IQ Running	6	Stepwise port of IQ code to GPUs based on established workflow followed for ICON-GPU implementation with OpenACC
modLSMcoup Running	6	Develop proof-of-concept for modular coupling of land surface and implement YAC coupler in ICON-eCLM coupling
CLEO Running	6	Coupling CLEO to ICON with YAC
PALM Waiting	6	Porting those of the PALM modules related to urban processes (especially radiation) to GPUs
MESSy-IMPORT Waiting	6	Revise the data import function of the Modular Earth Sub-model System (MESSy) for ICON/MESSy
MESSy-ComIn2 Waiting	6	ComIn integration time loop
PDAF2GPU Waiting	6	Porting PDAF to GPUs
PISM-AsyncIO Waiting	6	Resolve the issue with the I/O library for asynchronous output

+++ If you wish to unsubscribe from the natESM newsletter, please [click here](#). +++