



1. Generic questions/comments

Peter Braesicke: General observation:

Most ideas so far are really "big" - it is not clear to me, how a load balancing would work in a small team of four. Many projects (so far) could easily use all resources available ... is there a strategy: Many smaller requests should be favored over one large request? And another question: Are requests only possible if users want to use JSC or DKRZ machines? Will the team support portability to other systems as well?

- **Carsten Lemmen @Peter:** I agree, as we have lots of heterogeneity and wonderful diversity I think we could concentrate on interfaces, coupleability, modularity first, before doing the big ones. But these must be done as well "Leuchttürme" esp. those immediately deployable to exascale. ?
- **Hendryk Bockelmann @Peter:** good point, we should discuss this (either many smaller or some bigger requests) in the plenary after lunch ... about the HPC systems: it is not limited to JSC and DKRZ, especially if others systems are similar; but if a completely different system is not available to us (or we do not have the expertise yet), the support will be complicated

Daniel Klocke: A related general observation: most (all?) requests are from large research institutions, with big core funding, who could probably easily fund the developments from their own funds. Universities, which probably would need more help are missing. Is this ok, or is there a reason, that universities did not submit requests?

- **Hendryk Bockelmann @Daniel:** maybe it is related to the fact, that we required the main development to be long-term and with a large developer/user base?!
- **Michael Schulze @ Daniel:** Involvement of universities in the ESM strategy is definitely of great relevance for the success of the ESM strategy itself. Given the framework of the software sprints (long-term commitment by applicants to a certain ESM components) the imbalance between applications from universities and non-university institutions is not completely surprising (and reflecting the structure of the German science system). However, the sprints should generate added-value towards an ESM-system, which in turn should be beneficial to university users and should allow for projects that contribute to the further development of the natESM. The co-benefit of involving universities in the natESM was discussed in a DFG workshop on the role of universities in the strategy.
- **Carsten Lemmen @ Daniel:** Clear coupling interfaces and modularity (e.g. shown in ESMTTool today) would make it possible to support "small" university models that can integrate with the bigger model systems thru such interfaces.
- **Roland Potthast @ Daniel:** a lot of universities are involved, but focus on science, not infrastructure development. They ask the larger institutions to take that part, still being in the boat and using things immediately. This is a typical discussion we have at DWD and at MPI/DKRZ the same.



Oliver Bothe: Even then one may wonder why all requests (as far as I can see) are from within the HGF or from the MPI-M. Other institutions that do have models or use models from the Leibniz-Gemeinschaft (e.g., Tropos, IOW, IAP, PIK) or within the MPG (different from MPI-M) are missing. Also MARUM - which are university and have the long term support as far as I can tell - are missing.

- **Michael Schulze @Oliver:** At MARUM we are mainly using models, of which the code owners are outside Germany (e.g. CESM) - a prerequisite for the natESM.
- Oliver: Thanks, Michael, for both responses.

Ugur Cayoglu: Regarding the general process: Is it planned to have a rebuttal process?

Since the project proposals/descriptions should be rather concise and short, a rebuttal process might be a good idea. Especially since this is the first time such an opportunity exists. A rebuttal process would accelerate the process if there are some questions by the steering committee.

- **Michael Schulze @ Ugur:** The procedure is outlined in the call: "In case of ambiguities and follow-up questions during the technical assessment process, the applicant is invited to a video conference". This should be considered a clarification rather than a rebuttal.

Roland Potthast: Ownership, Governance and Sustainability: Taking software managed by larger international consortia is of course a point which needs careful discussion. As we know from many discussions in the wider ICON community it is very important to find efficient governance structures to be able to let developments find its way into sustainable branches of the models and their future development.

- **Carsten Lemmen @ Roland:** I saw many presentations with core development teams of 5-ish people, but rather large user bases. For such systems no governance structure is needed (as in ICON), but rather institutional commitment to support those few core devs.
- **Roland Potthast @ Carsten:** yes, I agree, if you own the software that's the important point C@R: "owning" is a difficult concept in open source. And the commonly used licenses GPL/Apache/BSD don't prevent people who disagree on governance (or disagree with the core devs) to do their own thing, even if the copyright is with one of our institutions.

Carsten Lemmen @Hendryk/SC: Will you ensure a transparent RF support process? For example, a ticketing system (possibly open)? Documentation of work done, such that others can benefit, not just reporting the outcome but the way.

- **Hendryk Bockelmann @ Carsten:** that is our goal; we will also schedule workshops to disseminate the best-practices



2. Specific questions for the applicants

2.1 Miguel Andrés-Martinez – AWI

Do you assume that the model codes are ported already? That can take quite some time.

- **Miguel Andrés – Martinez:** The models are currently running on Juwels, with some problems related to the Stages and modules that need to be tackled. Models will need to be ported from Mistral to Levante. We expect porting models from Mistral to Levante, to take from 2 to 3 weeks max, this is based in our experience porting models to other computers using ESM-Tools.

Carsten Lemmen: Down the road I'd like to discuss with you (maybe that could also be a sprint proposal) to add the ESMF/NUOPC coupler to the ESM Tools. Have you been discussing this already?

- **Miguel Andrés – Martinez:** Hi Carsten, we are always happy to include new models into our system. You can contact us sending an email to miguel.andres-martinez@awi.de, we will be very happy hearing from you

2.2 Dmitry Sidorenko – AWI

Hendryk Bockelmann: Why do you have to use a single binary for IFS-FESOM and not use classical coupling?

- **Dmitry Sidorenko:** There are several reasons:
 - (1) The single executable allows for sequential coupling without loss of computing resources (no waste of resources while idling).
 - (2) The coupled data assimilation that is used in IFS needs it to be a single executable.
 - (3) Single executable allows to immediately benefit from the recent developments in IFS. PS: We also make use of the coupler approach but do it with OpenIFS which is delayed in development from IFS.

2.3 Vera Fofonova – AWI

Sabine: How much would you already benefit from the FESOM porting?

- **Vera Fofonova:** FESOM-C is designed in a way to organize efficient coupling with FESOM2 (the same discretization - quasi-B, finite volume approach, the same manner of flux treatment, operates with unstructured meshes). However, FESOM-C is the different model with coastal functionality (different time stepping, numerical set for vertical dynamics realization, features as wetting/drying, turbulence closures and etc.).



Carsten Lemmen: is there any coupling infrastructure that you support in FESOM-C, and is this a shared infrastructure with FESOM-1/2? Is it ESMF/YAC/OASIS?

- **Vera Fofonova:** only one way coupling (from global to coastal) is realized with FESOM1/2. It shares infr. with global FESOM in a sense of general code organization.

2.4 Fatemeh Chegini – MPI-M

Sabine: How would you prioritize the 3 tasks you proposed?

2.5 Stephan Thober – UFZ

Julia Nabel: would this way of coupling later be an example for other impact modules

- **Daniel Klocke:** Yes, the relevant change for other models is, that, as Stephan explained, the use of the MPI space can be configured flexible, what is not possible right now. From this change, all models, which want to couple with ICON would profit.

Can coupling with YAC be considered in-situ coupling?

- I would say online coupling, maybe you mean the same
- **Panagiotis Adamidis:** YAC is working completely in parallel, in the sense that it can couple/interpolate fields, which are distributed over the local memories of the CPU's of a parallel machine.

The coupling activity should be coordinated between ICON-C (Consolidated) Commln Community Interface activity and YAC development used for coupling ICON-A with Waves, Ocean etc ...

- It is, but the community interface addresses much more.



2.6 Ali Hoshyaripour – KIT

Sabine: Which preparatory steps have you already taken? Are there "old" performance analyses available? Do you have a specific testing setup?

- **Ali Hoshyaripour:** We have performed an inventory of the modules but nothing more. There are no performance analysis but we intend to do it. And yes, we have tests with different complexity/ scopes

Have you ever run your code on JUWELS? How much time do you think it takes to port your model system to a new platform?

- **Ali Hoshyaripour:** Not yet. But it is similar to ICON: as ICON runs on JUWELS, I don't see any reason why ART should be different.

Domenico Taraborrelli: Which specific parts of ART do you suspect to take the most CPU time and memory?

- **Ali Hoshyaripour:** Good question: Chemistry and aerosol dynamics are the first suspects ;)

2.7 Peter Korn – MPI-M

Comment: This should be coordinated with the activities within DestinE climate adaptation twin as well as the ICON-C Community Interface activity.

Nikolay Koldunov: Do you plan to include CDO into the picture?

Johannes Quaas: this is very welcome for the user community, many thanks for the initiative!

2.8 Kerstin Hartung – DLR- PA

What is the main motivation to move to YAXT?

- **Astrid Kerkweg:** to make the communication of the MESSy models independent of the MPI communication within the individual base models. Only in this way,
 - MPI communication can be optimized for all MESSy base models at once
 - system specific optimizations can be applied much easier
 - this step is the basis to optimize the I/O facilities of MESSy, which is much more demanding with $O(1000)$ tracers instead of only $O(10)$ prognostic variables in GCMs



Carsten Lemmen: how does YAXT interface with other MPI abstractions/couplers (OASIS/ESMF/YAC)? and can it be used with those? Wouldn't they interfere, since these couplers also (can) take care of distribution and domain decompositions?

- **Panagiotis Adamidis:** Actually, YAC is using YAXT for some of the MPI communication. YAXT does not provide any interpolation scheme. It is handling only data communication and is suitable for complex communication patterns.
- **Moritz Hanke:** YAXT is an abstraction layer for MPI, which allows the easy implementation of complex collective communication patterns between processes (e.g. halo exchanges or transpositions). It does not interfere with coupling software.

2.9 Astrid Kerkweg – FZ Jülich + DLR-PA

Jochem Marotzke: How many of the challenges/tasks are due to MESSy itself (the coupling software) and how many tied to the component models? I guess the answer depends on how much time is spent in components versus coupling.

- **Astrid Kerkweg:** Actually none. In ICON (atmosphere) it is relatively fixed which processes are called after each other. And thus it is always clear (not dependent on the name list configuration) when the data needs to be copied from/to GPUs. The point is really, that each (additional) process (submodel in MESSy) might (will) be ported individually, however, without this functionality requested for we would have to copy for each process individually, which would end up in a lot of unnecessary, expensive copies (3x times order 1000 tracers) from one device to the other. Did I understand your question correct?

2.10 Joakim Kjellsson – GEOMAR

Jochem Marotzke: How would this relate to a national strategy. How much are you depending on decision taken by the NEMO developers?

- While most of the core development is done in the UK and France, GEOMAR is a core member of the DRAKKAR community which often guides the development of NEMO. For example, NEMO v4.2 will include a feature for regional vertical grid refinement which came out of requests from GEOMAR. Any suggestions we make such as loop re-ordering, would almost certainly make its way into the main NEMO trunk. This happened in 2019 when we, through another project, found a possible deadlock in the coupling with OASIS.



2.11 Daniel Caviedes Voullieme – FZ Jülich

Hendryk Bockelmann: Is there already a link to the Kokkos development team?

- **Daniel Caviedes Voullieme:** Yes, and they prove to be a quite accessible group.

2.12 Emil Stanev – HEREON

Lars: What is meant by 'operationalization', e.g. for PDAF and FABM?

- **Carsten Lemmen:** to be used not only for research but also in operational applications, think of HAB prediction, for example.

Lars: I'm wondering whether this could be achieved within a 4-week sprint given that neither DKRZ or JSC are operational centers. Actually, PDAF is already used operationally at BSH with HBM and CMEMS Baltic-MFC with NEMO - so there is experience.

- **Carsten Lemmen:** of course not, we are asking for support to find out about the requirements for implementing the items presented in a later sprint. We have PDAF-SCHISM, but not yet in operational settings. It might very quick to do this, but the assessment of the time required for implementation should be an outcome of the advisory phase.
- **Lars:** Thanks. I'm still wondering whether advice on operational implementations can be expected from DKRZ or JSC.
- **Carsten Lemmen:** we don't seek advice on the operationalization itself but on the software properties required for operationalization (think robustness, code quality, portability, interface completeness)
- **Lars:** Thanks for the clarification. Then I misunderstood Emil's slide and explanation on 'operationalization' in this context.

Oliver: In response to Jochem's question, it may be worth pointing to this schematic about the modelling strategy within the KS institute at Hereon

https://www.hereon.de/institutes/coastal_systems_analysis_modeling/research/gcoast/index.php.en of which Marcel showed another version.



2.13 Marcel Ricker – HEREON

Hendryk Bockelmann: if ~250 MPI-tasks is the current scaling limit; how will you achieve high-res?

- **Marcel Ricker:** The bottleneck seems to be sequential output. We think that the parallelization will increase the speed a lot. Furthermore, although 250 might sound not being much, wave models do not require as much resources as e.g. hydrodynamic models and so we are already able to conduct relatively long and/or highly resolved simulations.

Carsten Lemmen: regarding ownership the "problem" with open source is of course, that anyone can create forks and branch off. We are the owners of the original code and continue to support the GPL versioned open source WAM

- **Roland Potthast:** There is a WAM Version (ICON-WAVES) developed in cooperation with DWD ...
- **Marcel Ricken:** Correct. Mainly with the focus on the ICON coupling. And as Carsten wrote, the open source policy can induce problems when branches develop separately. We (Hereon) have to take of what goes into the original branch. Carsten: and with the GPL license we ensure that the downstream code developed at DWD can (legally) flow back into upstream dev.