

Dear Author,

Thank you for submitting your updated manuscript and for carefully replying to the reviewers' remarks. I consider that you have answered most of them but I still have some questions regarding some remarks of the first reviewer.

I also have some additional remarks that you will find here below.

Finally, while reading the manuscript, I could not resist making many propositions to improve the style of the text that you will also find starting on p.7 below. To construct this, I copied and pasted the content of your pdf manuscript in a Word file and made my changes in revision mode, so that you can easily identify them. The lines of the original manuscript appear separately, sorry for this but I think this was the most efficient way to propose many minor modifications. I hope you will find this useful to improve the style of the manuscript.

In the following, the reviewers remarks are in blue, your replies in black and my additional questions or comments in orange.

1) Reviewer 1

Title, Abstract

1) title Prepositions/articles missing. Also, I don't agree with "integrated" but suggest to use the word "modular". As the visualization has been proven, it is not a perspective "towards", but already production ready, so don't use "towards". I would rephrase the title (this is a mere suggestion) as "Co-processing online visualization: the Regional Earth System RegESM version 1.1".

In the current implementation, the adaptor code between NUOPC cap and ParaView, Catalyst is just a prototype that supports only for regional modelling cases with limited spatial grid support (structured and curvilinear grids). It still needs additional work to be more generic and adaptor code needs to be extended to include support for different mesh types such as unstructured grid for both global and regional applications. These are not implemented yet and needs to design adaptor code carefully for more generic applications. This will be done in the future.

Your argument about keeping "Toward" in the title is fine. However, you don't answer the reviewer's remark about the use of the word "integrated". I agree with the reviewer that use of "integrated" is not appropriate, as it seems contrary to the modularity of the approach. So what about "Toward modular in-situ visualization in Earth System Models: the regional modeling system RegESM 1.1".

2) abstract Fix language issues, such as: articles, avoid "being" before participles, check use of due to/because of, avoid overly use of "used", avoid "have to", do processes literally "move"? Avoid use of "basically", and many more. It is a major revision task to carefully check each sentence

Abstract section is modified based on your suggestions.

I still propose many language modifications in the attached ManuscriptAnnotated.pdf.

43) p8 l18f You name stability as an advantage of semi-implicit coupling. Also add drawbacks of semi-implicit and advantages of implicit time stepping.

The paragraph is extended to include more information about different coupling schemes.

To answer the referee's remark, you now mention the "implicit" coupling on l.30 but before defining it; the definition "The main difference between the implicit and semi-implicit coupling type is that the models interact on different time scales in implicit coupling scenarios. » comes only after but is not clear at all. Contrary to what was asked by the referee, I suggest to remove the text about the implicit coupling as it is not an option in RegESM (if I understand well).

46) p8 l31f The description of regridding could be a subsection, section 2.6 "Connector". How do other coupled model systems handle this (any references available?). Recently, I implemented into MOSSCO a very similar scheme using bilinear for interpolation and nearest neighbour for extrapolation, so I believe this is a very good way to handle this (but it is not published yet). Possibly, describe how ESMF handles this rather than pointing vaguely to ongoing development.

I prefer to keep it as it is. The special interpolation implemented in RegESM modeling system is a part of driver.

It is OK for me not to describe how ESMF handles this. But your description on p.9, l.19-32 is very hard to follow.

In particular, I don't understand at all what "According to the algorithm, the mapped grid points have same land-sea mask type in both model components (i.e., both are sea or land). On the other hand, the land-sea mask type does not match completely in the case of unmapped grid points" means.

Can you try to review and simplify it, maybe giving the essence of the method and not all the technical steps? Also on Fig. 4, a "on" is missing between "only" and "grid" in "All interpolations are performed only grid points over SEA".

47) p9 l14 Align your description with the terms used in Figure 4 (see comments below on this figure). We used field names (Field_A, Field_B and Field_C) in also manuscript. The paragraph is also extended to describe the algorithm clearly.

I am sorry to say that I think that the added paragraph does not help understanding the algorithm.

49) p9 l28 Define "conventional" or give a reference to such approaches. Define pipeline in this context.

The conventional co-processing enabled simulation systems interacts with single physical model component such as atmosphere along with co-processing support. This is defined also in text and references are added.

... A visualization pipeline integrates a data flow network in which computation is described as a collection of executable modules that are connected in a directed graph representing how data moves between modules (Moreland, 2013). There are three types of modules: sources (file readers and synthetic data generators), filters (transforms data), and sinks (file writers and rendering module that provide images to a user interface) in the visualization pipeline. ...

There is some incoherency in the first paragraph of section 3. The paragraph starts by

describing the “conventional co-processing” and without transition discusses the NUOPC cap (which, if I understand well, is part of the novel approach and not the conventional one); please clarify.

Also, in the abstract and introduction, you use the word “conventional” associated to “post-processing”, which may lead to some confusion. I would advice to use “traditional” instead of “conventional” when qualifying the post-processing and keep “conventional” when qualifying the co-processing.

52) p14 l26 The “driver component introduces additional 5-10%”. In ESMF documentations, usually below 5% overhead is assumed. this makes sense as you to the SMM twice for each interpolation of mask extrapolation. It would be interesting to see whether the generic implementation of extrapolation is more efficient in ESMF 8 (out of scope of this MS).

I agree. As far as i know, the ESMF developers was trying to find a way to combine weight matrix of multiple interpolation (including extrapolation) into a single one. Similar methodology is using in transformation (scale, rotation, translation etc.) matrixes using homogeneous coordinates but i am not sure it could be used in regriding or not.

Can you specify how you calculate the overhead? Do you compare the CPL wall clock-time to the sum of the standalone OCN and ATM wall clock time as they run sequentially? Please clarify.

53) p14 l35 “ includes vertical interpolation to map data from sigma coordinates to height coordinates”. Why is this interpolation step necessary? In principle, the visualization should equally work on sigma layers. And is the interpolation performed within the ESMF connector, or within a Paraview pipeline?

I did not try yet but yes the visualization works with sigma coordinates without any problem. The vertical interpolation is introduced to have a consistency in the data coming from atmosphere and ocean. By this way, the vertical scales of the data can be compared without any problem. In general, both atmosphere and ocean components use sigma coordinates with different definitions. The interpolation is done in ESMF cap not in ParaView side. It could be done in ParaView side also but in this case all grid related parameters such as stretching functions etc. need to be passed to ParaView, Catalyst and special ProgrammableFilter need to be developed.

Please add the justification of the vertical interpolation in the manuscript.

54) p15 l2 the computational demand for visualization “require 10-40% extra” is rather high and its evaluation would probably change depending on whether it is more 10% (acceptable) or 40% (not acceptable). The benchmark using software rendering with Mesa is not ideal to demonstrate this, especially as many new HPC systems come with dedicated GPU nodes that could much reduce the visualization overhead.

It is clear that Mesa is not right way to measure the performance of the co-processing component. The overall performance of the used computing system dominates the results in our case but we don't have accesses to a system with GPU accelerators. In the early stage of the performance benchmark, Dell provided us a grant in newly installed University Manchester system (10000 core/hour) but the system is mainly designed for GPU intensive and AI type researches and the nodes have small number of core (10 cores) but four NVIDIA P100 GPU. In fact, our test simulations are CPU intensive and

requires high amount of compute resource (or cores) for the high-resolution simulations and rendering only done on a single GPU (limitation of ParaView, Catalyst and used filters). As a result, couple of benchmark run spend all allocated resources (unused GPU resources was also accounted) and the system was not used for the benchmark. There was also issue related with buggy NVIDIA driver and ParaView protobuf library. In the future, along with the development of m-VTK (GPU accelerated filters) it could be possible to use multi-GPU systems.

I agree with the reviewer that 40% is very significant. I suppose it is linked to the fact that the components are all run sequentially (if I understand well). You should say something about the overhead if the co-processing was run concurrently; more processes would be needed but the wall-time would probably not increase.

59) Figure 4 Rather unreadable. Please redesign. (1) Avoid bridges (this is possible by routing the arrow "Create ESMF FIELD_B" north to "Create ESMF FIELD_C". (2) don't use same line thickness/color for outlining the boxes around src/dst grid, rather make them visually less intrusive to highlight the flowchart. (3) Abstract away 1.0, LARGEST, "over sea" or other value-based masking, if possible.

The figure is modified. The hop is solved by changing the position of the arrow. The dashed line is used in the boxes and color is changed to grey to move focus to workflow. It is hard to abstract used values and mask values. The algorithm is tested with this configuration and it is better to keep them as it is to have a reproducible result (at least for finding mapped and unmapped grid points). I am also using same methodology using NCL ESMF interface to create forcing for ocean models and it works better than standard extrapolation method (based on solution of Possion Eq.) provided by NCL itself.

I cannot really comment the figures, as I don't understand the algorithm, see my remark above.

61) Figures 9/10 It is very difficult to compare ATM versus OCN performance using different y scales. I recommend to use the same ylog-scale for all panels in figs 9 and 10. Also try loglog. Tile sizes are hard to read, please move text from lines. Calculate all speedup relative to 140 cores, so that you don't have different reference points. Explain model versions like RegCM_r6274, delete UHeM. Avoid overlapping graphics and text.

The speed-up calculation of coupled model simulations starts from 140 core because this is the minimum number of cores that allows to run the modeling system. For standalone simulations, it is possible to run the model with 28 cores (single node). As a result, it was used 28 cores for standalone and 140 cores for coupled model simulations to calculate the speedup. The plots are modified to have consistent x and y axis (log scale) as well as number of cores used in the speed-up calculation (140 cores) except plot for standalone ocean model. In this case, using same log scale for it does not give good results as expected for example the effect of tile configuration cannot be seen clearly because the benchmark results fall into a very narrow interval. To that end, i prefer to keep results of standalone ocean simulations as it is. I also increased font sizes and move the text little bit far from the markers for standalone ocean model plot. The texts that indicate model version and name of the computing system that is used in the benchmark are removed because they are already mentioned in the manuscript.

You should mention somewhere in the text that figures 9 and 10 show the wall clock time and the speed-up. Please define precisely how you calculate the speed-up either in the text or in the captions.

Additional remarks

- For OASIS3-MCT, please cite: “Craig A., Valcke S., Coquart L., 2017: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, Geoscientific Model Development, 10, pp. 3297-3308, doi:10.5194/gmd-10-3297-2017”
- Many places in the text, you use « ParaView, Catalyst » to design the ParaView co-processing plugin. Please use « ParaView Catalyst » without the comma or « ParaView/Catalyst »
- COP is used sometimes to design the co-processing component and sometimes the three-component coupled system simulations. This is confusing. Please use “co-processing component” for the component and keep COP to design only the three-component coupled system simulations. Please define COP and CPL the first time it appears in the text.
- Fig. 7 is misleading. It looks like the ATM and OCN components are coupled **through** the co-processing component. If I understand well, this figure should just illustrate the interaction between ATM and the co-processing component on one side, and the interaction between OCN and the co-processing component on the other side. If I am right could it be possible to split the figure into two parts so to avoid the confusion?
- In section 4.2, an HR (inner) atmosphere is nested in an LR (outer) atmosphere. In section 4.3, HR and LR are used, if I understand well, for two different atmospheres covering the whole atmospheric domain (no nesting); this is confusing. Can you keep HR and LR for section 4.3 only, and change the wording when referring to the HR-inner atmosphere nested in the LR-outer atmosphere in section 4.2?
- P.14 , l.4: “their THREDDS server”; can you explain what THREDDS means or give a link?
- P.14, l.30: You refer to section 2.5 for details on the limitation of the co-processing about its sequential type execution, but I don’t see where this is detailed in section 2.5. Can you clarify or point me to the exact paragraph?
- Table 1: Pipeline details are not readable. Table 1 caption: please remove “in” in “... are shown in here ...”
- P.15, l.6-8: I don’t understand why you write “It is also shown that around 588 processors, which is the highest available compute resource, the communication among the processors dominate the benchmark results and even HR case does not gain further performance”: the HR curve (triangles) does not flatten as the LR curve. Please clarify.

- Fig. 9: I don't understand what the envelope represents and what do you mean by "as a line". Why isn't the best configuration the lower limit of the envelope?
- P. 15, l.23: I don't understand what "acceptable when increased number of MPI communication between the components are considered" means in this context. Is it a justification/explanation of the 5-10% overhead? Please rephrase.
- I think you don't need to put a capital letter to each word in titles in English.
- Color scales in Fig 12, 13 and 14 are not readable.
- P.19, l.15: Please rephrase "have higher water content in a decreasing trend with height and spatial distribution" as it does not seem grammatically correct to me.
- P.21, l.23: Please modify the sentence "that mounts the user mode components of the driver and the GPUs into the container at lunch" as it is not understandable for a non-expert reader.

Towards modular in-situ visualization in Earth System Models: the regional modeling system RegESM 1.1

Ufuk Utku Turuncoglu

Informatics Institute, Istanbul Technical University, 34469, Istanbul, Turkey
Correspondence: Ufuk Utku Turuncoglu (ufuk.turuncoglu@itu.edu.tr)

Abstract. The data volume produced by regional and global multi-component Earth System Models is rapidly increasing because of the improved spatial and temporal resolution of the model components, and the sophistication of the numerical models regarding represented physical processes and their complex non-linear interactions. In particular, very small time steps need to be defined in non-hydrostatic high-resolution modeling applications to represent the evolution of the fast-moving processes such as turbulence, extra-tropical cyclones, convective lines, jet streams, internal waves, vertical turbulent mixing and surface gravity waves. Consequently, the employed small time steps cause extra computation and disk input-output overhead in the modeling system even if today's most powerful high-performance computing and data storage systems are considered. Analysis of the high volume of data from multiple Earth System Model components at different temporal and spatial resolution also poses a challenging problem to efficiently perform integrated data analysis of the massive amounts of data when relying on the conventional post-processing methods. This study mainly aims to explore the feasibility and added value of integrating existing in-situ visualization and data analysis methods within the model coupling framework. The objective is to increase interoperability between Earth System multi-component code and data processing systems by providing an easy-to-use, efficient, generic and standardized modeling environment. The new data analysis approach enables simultaneous analysis of the vast amount of data produced by multi-component regional Earth System Models during the runtime. The presented methodology also aims to create an integrated modeling environment for analyzing fast-moving processes and their evolution both in time and space to support a better understanding of the underlying physical mechanisms. The state-of-art approach can also be employed to solve common problems in the model development cycle, e.g. designing a new sub-grid scale parameterization that requires inspecting the integrated model behavior at a higher temporal and spatial scale simultaneously and supporting visual debugging of the multi-component modeling systems, which usually are not facilitated by existing model coupling libraries and modeling systems.

1 Introduction

The multi-scale and inherently coupled Earth System Models (ESMs) make them challenging to study and understand. Rapid developments in Earth system science, as well as in high-performance computing and data storage systems, have enabled fully coupled regional or global ESMs to better represent relevant processes, complex climate feedbacks, and interactions among the coupled components. In this context, regional ESMs are employed when the spatial and temporal resolution of the global climate models are not sufficient to resolve local features such as complex topography, land-sea gradients and the influence of human activities in a smaller spatial scale. Along with the development of the modeling systems, specialized software libraries for the model coupling become more and more critical to reduce the complexity of the coupled model development and increase the interoperability, reusability, and efficiency of the existing modeling systems. Currently, the existing model coupling software libraries have two main categories: couplers and coupling frameworks.

Couplers are mainly specialized in performing specific operations more efficiently and quickly such as coordination of components and interpolation among model components. For example, OASIS3 (Valcke, 2013) uses multiple executable approaches for coupling model components but sequentially performing internal algorithms such as sparse matrix multiplication (SMM) operation for interpolation among model grids become a bottleneck along with increased spatial resolution of the model components. To overcome the problem, OASIS4 uses parallelism in its internal algorithms (Redler et al., 2010), and OASIS3-MCT interfaced with the Model Coupling Toolkit (MCT; Jacob et al., 2005; Larson et al., 2005) provides a parallel implementation of interpolation and data exchange. Besides generic couplers like OASIS, domain specific couplers such as Oceanographic Multi-purpose Software Environment (OMUSE; Peluassy et al., 2017) that aims to provide a homogeneous environment for ocean modeling to make verification of simulation models with different codes and numerical methods and Community Surface Dynamics Modeling System (CSDMS; Overeem et al., 2013) to develop integrated software modules for modeling of Earth surface processes are introduced.

A coupling framework is an environment for coupling model components through a standardized calling interface and aims to reduce the complexity of regular tasks such as performing spatial interpolation across different computational grids and transferring data among model components to increase the efficiency and

Sophie Valcke 21/11/y 18:28

Supprimé: integrated

Sophie Valcke 21/11/y 18:28

Supprimé: regional modeling system

Sophie Valcke 21/11/y 18:32

Supprimé: are

Sophie Valcke 21/11/y 18:36

Supprimé: used

Sophie Valcke 21/11/y 18:38

Supprimé: s

Sophie Valcke 21/11/y 18:39

Supprimé: by

Sophie Valcke 21/11/y 18:39

Supprimé: available today

Sophie Valcke 21/11/y 18:41

Supprimé: simulation

Sophie Valcke 21/11/y 18:42

Supprimé:

Sophie Valcke 21/11/y 18:42

Supprimé:

Sophie Valcke 21/11/y 18:42

Supprimé: for Earth system science applications

Sophie Valcke 21/11/y 18:43

Supprimé: in

Sophie Valcke 21/11/y 18:45

Supprimé: :

Sophie Valcke 21/11/y 18:43

Supprimé: parametrizations

Sophie Valcke 21/11/y 18:44

Supprimé: in

Sophie Valcke 21/11/y 19:02

Supprimé: but

Sophie Valcke 21/11/y 19:02

Supprimé: coupler

Sophie Valcke 21/11/y 19:02

Supprimé: to develop OASIS3-MCT that

Sophie Valcke 23/11/y 11:06

Supprimé: as

interoperability of multi-component modeling systems. Besides, the synchronization of the execution of individual model components, a coupling framework can simplify the exchange of metadata related to model components and exchanged fields through the use of existing conventions such as CF (Climate and Forecast) convention. The Earth System Modeling Framework (ESMF) is one of the most famous examples of this approach (Theurich et al., 2016). The ESMF consists of a standardized superstructure for coupling components of Earth system applications through a robust infrastructure of high-performance utilities and data structures that ensure consistent component behavior (Hill et al., 2004). The ESMF framework is also extended to include the National Unified Operational Prediction Capability (NUOPC) layer. The NUOPC layer simplifies component synchronization and run sequence by providing additional programming interface between coupled model and ESMF framework through the use of a NUOPC “cap”. In this case, a NUOPC “cap” is a Fortran module that serves as the interface to a model when it is used in a NUOPC-based coupled system. The term “cap” is used because it is a small software layer that sits on top of a model code, making calls into it and exposing model data structures in a standard way. In addition to generic modeling framework like ESMF, Modular System for Shelves and Coasts (MOSSCO; Lemmen et al., 2018) creates a state-of-art domain and process coupling system by taking advantage of both ESMF and Framework for Aquatic Biogeochemical Models (FABM; Bruggeman and Bolding, 2014) for marine coastal Earth system community. The recent study of Alexander and Easterbrook (2015) to investigate the degree of modularity and design of the existing global climate models reveals that the majority of the models use central couplers to support data exchange, spatial interpolation, and synchronization among model components. In this approach, direct interaction does not have to occur between individual model components or modules, since the specific coupler component manages the data transfer. This approach is also known as the hub-and-spoke method of building a multi-component coupled model. A key benefit of using a hub-and-spoke approach is that it creates a more flexible and efficient environment for designing sophisticated multi-component modeling system regarding represented physical processes and their interactions. The development of the more complex and high-resolution modeling systems leads to an increased demand for both computational and data storage resources. In general, the high volume of data produced by the numerical modeling systems may not allow storing all the critical and valuable information to use later, despite recent advances in storage systems. As a result, the simulation results are stored in a limited temporal resolution (i.e., monthly averages), which are processed after numerical simulations finished (post-processing). The poor representation of the results of numerical model simulations prevents to analyze the fast-moving processes such as extreme precipitation events, convection, turbulence and non-linear interactions among the model components in a high temporal and spatial scale with the conventional post-processing approach. The analysis of leading high-performance computing systems reveals that the rate of disk input-output (I/O) performance is not growing at the same speed as the peak computational power of the systems (Ahern, 2012; Ahrens, 2015). The recent report of U.S. Department of Energy (DOE) also indicates that the expected rate of increase in I/O bandwidth (100 times) will be slower than the peak system performance (500 times) of the new generations of exascale computers (Ashby et al., 2010). Besides, the movement of large volumes of data across relatively slow network bandwidth servers fails to match the ultimate demands of data processing and to archive tasks of the present high-resolution multi-component ESMs. As a result, the conventional post-processing approach has become a bottleneck in monitoring and analysis of fast-moving processes that require very high spatial resolution, due to the present technological limitations in high-performance computing and storage systems (Ahrens et al., 2014). In the upcoming computing era, state-of-art new data analysis and visualization methods are needed to overcome the above limitations evocatively. Besides the conventional data analysis approach, the so-called in-situ visualization and co-processing approaches allow researchers to analyze the output while running the numerical simulations simultaneously. The coupling of computation and data analysis helps to facilitate efficient and optimized data analysis and visualization pipelines and boosts the data analysis workflow. Recently, a number of in-situ visualization systems for analyzing numerical simulations of Earth system processes have been implemented. For instance, the ocean component of Model for Prediction Across Scales (MPAS) has been integrated with an image-based in-situ visualization tool to examine the critical elements of the simulations and reduce the data needed to preserve those elements by creating a flexible work environment for data analysis and visualization (Ahrens et al., 2014; O’Leary et al., 2016). Additionally, the same modeling system (MPAS-Ocean) has been used to study eddies in large-scale, high-resolution simulations. In this case, the in-situ visualization workflow is designed to perform eddy analysis at higher spatial and temporal resolutions than available with conventional post-processing facing storage size and I/O bandwidth constraints (Woodring et al., 2016). Moreover, a regional weather forecast model (Weather Research and Forecasting Model; WRF) has been integrated with in-situ visualization tool to track cyclones based on an adaptive algorithm (Malakar et al., 2012). Despite the lack of generic and standardized implementation for integrating model components with in-situ visualization tools, the previous studies have shown that in-situ visualization can produce analyses of simulation results, revealing many details in an efficient and optimized way. It is evident that more generic implementations could facilitate smooth integration of the

existing standalone and coupled ESMs with available in-situ visualization tools (Ahrens et al., 2005; Ayachit, 2015; Childs et al., 2012) and improve interoperability between such tools and non-standardized numerical simulation codes.

The main aim of this paper is to explore the added value of integrating in-situ analysis and visualization methods with a model coupling framework (ESMF) to provide in-situ visualization for easy to use, generic, standardized and robust scientific applications of Earth system modeling. The implementation allows existing ESMs coupled with the ESMF library to take advantage of in-situ visualization capabilities without extensive code restructuring and development. Moreover, the integrated model coupling environment allows sophisticated analysis and visualization pipelines by combining information coming from multiple ESM components (i.e., atmosphere, ocean, wave, land-surface) in various spatial and temporal resolutions. Detailed studies of fundamental physical processes and interactions among model components are vital to the understanding of complex physical processes and could potentially open up new possibilities for the development of ESMs.

The Design of the Modeling System

The RegESM (Regional Earth System Model; 1.1) modeling system can use five different model components to support many different modeling applications that might require detailed representation of the interactions among different Earth system processes (Fig. 1a-b). The implementation of the modeling system follows the hub-and-spoke architecture. The driver that is responsible for the orchestration of the overall modeling system resides in the middle and acts as a translator among model components (atmosphere, ocean, wave, river routing, and co-processing). In this case, each model component introduces its NUOPC cap to plug into the modeling system. The modeling system is validated in different model domains such as Caspian Sea (Turuncoglu et al., 2013), Mediterranean Basin (Surenkok and Turuncoglu, 2015; Turuncoglu and Sannino, 2017), and Black Sea Basin.

2.1 Atmosphere Models (ATM)

The flexible design of RegESM modeling system allows choosing a different atmospheric model component (ATM) in the configuration of the coupled model for a various type of application. Currently, two different atmospheric model is compatible with RegESM modeling system: (1) RegCM4 (Giorgi et al., 2012), which is developed by the Abdus Salam International Centre for Theoretical Physics (ICTP) and (2) the Advanced Research Weather Research and Forecasting (WRF) Model (ARW; Skamarock et al., 2005), which is developed and sourced from National Center for Atmospheric Research (NCAR). In this study, RegCM 4.6 is selected as an atmospheric model component because the current implementation of WRF coupling interface is still experimental and does not support coupling with co-processing component yet, but the next version of the modeling system (RegESM 1.2) will be able to couple WRF atmospheric model with co-processing component. The NUOPC cap of atmospheric model components defines state variables (i.e., sea surface temperature, surface wind components), rotates the winds relative to Earth, apply unit conversions and perform vertical interpolation to interact with the newly introduced co-processing component.

2.1.1 RegCM

The dynamical core of the RegCM4 is based on the primitive equation, hydrostatic version of the National Centre for Atmospheric Research (NCAR) and Pennsylvania State University mesoscale model MM5 (Grell, 1995). The latest version of the model (RegCM 4.6) also supports non-hydrostatic dynamical core to support applications with high spatial resolutions (< 10 km). The model includes two different land surface models: (1) Biosphere-Atmosphere Transfer Scheme (BATS; Dickinson et al., 1989) and (2) Community Land Model (CLM), version 4.5 (Tawfik and Steiner, 2011). The model also includes specific physical parameterizations to define air-sea interaction over the sea and lake (one-dimensional lake model; Hostetler et al., 1993). The Zeng Ocean Air-Sea Parameterization (Zeng et al., 1998) is extended to introduce the atmosphere model as a component of the coupled modeling system. In this way, the atmospheric model can exchange both two and three dimensional fields with other model components such as an ocean, wave and river routing components that are active in an area inside of the atmospheric model domain as well as in-situ visualization component.

2.1.2 WRF

The WRF model consists of fully compressible non-hydrostatic equations, and the prognostic variables include the three dimensional wind, perturbation quantities of pressure, potential temperature, geo-potential, surface pressure, turbulent kinetic energy and scalars (i.e., water vapor mixing ratio, cloud water). The model is suitable for a broad range of applications and has a variety of options to choose parameterization schemes for the

planetary boundary layer (PBL), convection, explicit moisture, radiation, and soil processes to support analysis of different Earth system processes. The PBL scheme of the model has a significant impact on exchanging moisture, momentum, and energy between air and sea (and land) due to the used alternative surface layer options (i.e., drag coefficients) in the model configuration. A few modifications are done in WRF (version 3.8.1) model itself to couple it with RegESM modeling system. These modifications include rearranging of WRF time-related subroutines, which are inherited from the older version of ESMF Time Manager API (Application Programming Interface) that was available in 2009, to compile model with the newer version of ESMF library (version 7.1.0) together with the older version that requires mapping of time manager data types between old and new versions.

2.2 Ocean Models (OCN)

The current version of the coupled modeling system supports two different ocean model components (OCN): (1) Regional Ocean Modeling System (ROMS revision 809; Shchepetkin and McWilliams, 2005; Haidvogel et al., 2008), which is developed and distributed by Rutgers University and (2) MIT General Circulation Model (MITgcm version c63s; Marshall et al., 1997a, b). In this case, ROMS and MITgcm models are selected due to their large user communities and different vertical grid representations. Although the selection of ocean model components depends on user experience and application, often the choice of vertical grid system has a determining role in some specific applications. For example, the ROMS ocean model uses terrain following (namely s-coordinates) vertical grid system that allows a better representation of the coastal processes but MITgcm uses z levels generally used for applications that involve open oceans and seas. Similar to the atmospheric model component, both ocean models are slightly modified to allow data exchange with the other model components. In the current version of the coupled modeling system, there is no interaction between wave and ocean model components, which could be crucial for some applications (i.e., surface ocean circulation and wave interaction) that need to consider the two-way interaction between waves and ocean currents. The exchange fields defined in the coupled modeling system between ocean and atmosphere strictly depend on the application and the studied problem. In some studies, the ocean model requires heat, freshwater and momentum fluxes to be provided by the atmospheric component, while in others, the ocean component retrieves surface atmospheric conditions (i.e., surface temperature, humidity, surface pressure, wind components, precipitation) to calculate fluxes internally, by using bulk formulas (Turuncoglu et al., 2013). In the current design of the coupled modeling system, the driver allows selecting the desired exchange fields from the predefined list of the available fields. The exchange field list is a simple database with all fields that can be exported or imported by the component. In this way, the coupled modeling system can be adapted to different applications without any code customizations in both the driver and individual model components.

2.2.1 ROMS

The ROMS is a three-dimensional, free-surface, terrain-following numerical ocean model that solves the Reynolds-averaged Navier-Stokes equations using the hydrostatic and Boussinesq assumptions. The governing equations are in flux form, and the model uses Cartesian horizontal coordinates and sigma vertical coordinates with three different stretching functions. The model also supports second, third and fourth order horizontal and vertical advection schemes for momentum and tracers via its preprocessor flags.

2.2.2 MITgcm

The MIT general circulation model (MITgcm) is a generic and widely used ocean model that solves the Boussinesq form of Navier-Stokes equations for an incompressible fluid. It supports both hydrostatic and non-hydrostatic

Sophie Valcke 22/11/y 10:17

Supprimé: and

Sophie Valcke 22/11/y 10:17

Supprimé: the

Sophie Valcke 23/11/y 11:39

Supprimé: known

applications with a spatial finite-volume discretization on a curvilinear computational grid. The model has an implicit free surface in the surface and partial step topography formulation to define vertical depth layers. The MITgcm model supports different advection schemes for momentum and tracers such as centered second order, third-order upwind and second-order flux limiters to support a variety of applications. The model used in the coupled modeling system was slightly modified by ENEA to allow data exchange with other model components. The detailed information about the regional applications of the MITgcm ocean model is described in the study of Artale et al. (2010) using PROTHEUS modeling system specifically developed for the Mediterranean Sea.

2.3 Wave Model (WAV)

Surface waves play a crucial role in the dynamics of PBL in the atmosphere and the currents in the ocean. Therefore, the wave component is included in the coupled modeling system to have a better representation of atmospheric PBL and surface conditions (i.e., surface roughness, friction velocity, wind speed). In this case, the wave component is based on WAM Cycle-4 (4.5.3-MPI). The WAM is a third-generation model without any assumption on the spectral shape (Monbaliu et al., 2000). It considers all the main processes that control the evolution of a wave field in deep water, namely the generation by wind, the nonlinear wave-wave interactions, and also white-capping. The model was initially developed by Helmholtz-ZentrumGeesthacht (GKSS, now HZG) in Germany. The original version of the WAM model was slightly modified to retrieve surface atmospheric conditions (i.e., wind speed components or friction velocity and wind direction) from the RegCM4 atmospheric model and to send back calculated surface roughness. In the current version of the modeling system, wave component cannot be coupled with the WRF model due to the missing modifications in the WRF side. In the RegCM4, the received surface roughness is used to calculate air-sea transfer coefficients and fluxes over sea using Zeng ocean air-sea parameterization (Zeng et al., 1998). In this design, it is also possible to define a threshold for maximum roughness length (the default value is 0.02 m) and friction velocity (the default value is 0.02 m) in the configuration file of RegCM4 to ensure the stability of the overall modeling system. Initial investigation of the added value of atmosphere-wave coupling in the Mediterranean Sea can be found in Surenkok and Turuncoglu (2015).

2.4 River Routing Model (RTM)

To simulate the lateral freshwater fluxes (river discharges) at the land surface and to provide river discharge to ocean model component, the RegESM modeling system uses Hydrological Discharge (HD, version 1.0.2) model developed by Max Planck Institute (Hagemann and Dumenil, 1998; Hagemann and Lydia, 2001). The model is designed to run in a fixed global regular grid with 0.5° horizontal resolution using daily time series of surface runoff and drainage as input fields. In that case, the model uses the pre-computed river channel network to simulate the horizontal transport of the runoff within model watersheds using different flow processes such as overland flow, baseflow and riverflow. The river routing model (RTM) plays an essential role in the freshwater budget of the ocean model by closing the water cycle between the atmosphere and ocean model components.

Sophie Valcke 22/11/y 10:18

Supprimé: is

Sophie Valcke 22/11/y 10:18

Supprimé: was initially

Sophie Valcke 22/11/y 10:19

Supprimé: is

Sophie Valcke 22/11/y 10:19

Supprimé: is

Sophie Valcke 22/11/y 10:20

Supprimé: The i

Sophie Valcke 22/11/y 10:20

Supprimé: results to

Sophie Valcke 22/11/y 10:20

Supprimé: e

Sophie Valcke 22/11/y 10:22

Commentaire [1]: To distinguish from the « In this case » below which refers to the current setting, if I understand well.

Sophie Valcke 22/11/y 10:21

Supprimé: this

The original version of the model was slightly modified to support interaction with the coupled model components. To close water cycle between land and ocean, model retrieves surface and sub-surface runoff from the atmospheric component (RegCM or WRF) and provides estimated river discharge to the selected ocean model component (ROMS or MITgcm). In the current design of the driver, rivers can be represented in two different ways: (1) individual point sources that are vertically distributed to model layers, and (2) imposed as freshwater surface boundary condition like precipitation (P) or evaporation minus precipitation (E-P). In this case, the driver configuration file is used to select the river representation type (1 or 2) for each river individually. The first option is preferred if river plumes need to be defined correctly by distributing river discharge vertically among the ocean model vertical layers. The second option is used to distribute river discharge to the ocean surface when there is a need to apply river discharge to a large areal extent close to the river mouth. In this case, a special algorithm implemented in NUOPC cap of ocean model components (ROMS and MITgcm) is used to find affected ocean model grids based on the effective radius (in km) defined in the configuration file of the driver.

2.5 The Driver: RegESM

The RegESM (version 1.1) is completely redesigned and improved version of the previously used and validated coupled atmosphere-ocean model (RegCM-ROMS) to study the regional climate of Caspian Sea and its catchment area (Turuncoglu et al., 2013). To simplify the design and to create more generic, extensible and flexible modeling system that aims to support easy integration of multiple model components and applications, the RegESM uses a driver to implement the hub-and-spoke approach.

In this case, all the model components are combined using ESMF (version 7.1.0) framework to structure coupled modeling system. The ESMF framework is selected because of its unique online re-gridding capability, which allows the driver to perform different interpolation types (i.e., bilinear, conservative) over the exchange fields (i.e., sea surface temperature, heat and momentum fluxes) and the NUOPC layer. The NUOPC layer is a software layer built on top of the ESMF. It refines the

capabilities of ESMF by providing a more precise definition of a component model and how components should interact and share data in a coupled system. The ESMF also provides the capability of transferring computational

grids in the model component memory, which has critical importance in the integration of the modeling system with

a co-processing environment (see also Sect. 3). The RegESM modeling system also uses ESMF and NUOPC layer to support various configuration of component interactions such as defining multiple coupling time steps among the model components.

An example configuration of the four-component (ATM, OCN, RTM, and WAV) coupled modeling system can be seen in

Fig. 2. In this case, the RTM component runs in a daily time step (slow) and interacts with ATM and OCN components, but

ATM and OCN components can interact each other more frequently (fast) such as every three hours.

The interaction (also called as run sequences) among the model components and driver are facilitated by the connector

components provided by NUOPC layer. Connector components are mainly used to create a link between individual model components and driver. In this case, the number of active components and their interaction determines the number of connector

component created in the modeling system. The interaction between model components can be in two way: (1) bi-directional

such as atmosphere and ocean coupled modeling system or (2) unidirectional such as atmosphere and co-

Sophie Valcke 22/11/y 10:21

Supprimé: is

Sophie Valcke 22/11/y 10:23

Supprimé: what it means for a model to be

Sophie Valcke 22/11/y 10:23

Supprimé: memory among

Sophie Valcke 22/11/y 10:23

Supprimé: s

processing modeling

system. In the uni-directional case, the co-processing component does not interact with the atmosphere model and only process

retrieved information; thus there is one connector component.

The RegESM modeling system can use two different types of time-integration coupling scheme between the atmosphere and ocean components: (1) explicit and (2) semi-implicit (or leap-frog) (Fig. 3). In the explicit type coupling, two

connector components (ATM-OCN and OCN-ATM direction) are executed concurrently at every coupling time step and model components

start and stop at the same model time (Fig. 3a). In the semi-implicit coupling type (Fig. 3b), the ocean model receives surface boundary conditions from the atmospheric model at one coupling time step ahead of the current ocean model time.

The implicit and semi-implicit coupling aimed at lowering the overall computational cost of a simulation by increasing stability for longer coupling time steps. The main difference between the implicit and semi-implicit coupling type is that the models interact on different time scales in implicit coupling scenarios.

As described earlier, the execution of the model components is controlled by the driver. Both sequential and concurrent

execution of the model components is allowed in the current version of the modeling system. If the model components and

the driver are configured to run in sequence on the same set of PETs (Persistent Execution Threads), then the modeling

system executes in a sequential mode. This mode is a much more efficient way to run the modeling system in case of limited

computing resources. In the concurrent type of execution, the model components run in mutually exclusive sets of PETs, but the

NUOPC connector component uses a union of available computational resources (or PETs) of interacted model components.

By this way, the modeling system can support a variety of computing systems ranging from local servers to large computing systems that could include high-speed performance networks, accelerators (i.e., Graphics Processing Unit or GPU) and parallel

I/O capabilities. The main drawback of concurrent execution approach is to assign correct amount of computing resource to

individual model components, which is not an easy task and might require an extensive performance benchmark of a specific

configuration of the model components, to achieve best available computational performance. In this case, a load-balancing

analysis of individual components and driver play a critical role in the performance of the overall modeling system. For example, the LUCIA (Load-balancing Utility and Coupling Implementation Appraisal) tool can be used to collect all required

information such as waiting time and calculation time of each system components for a load-balancing analysis in the OASIS3-MCT based coupled system.

In general, the design and development of the coupled modeling systems involve a set of technical difficulties that arise due to the usage of the different computational grids in the model components. One of the most common examples is the mismatch between the land-sea masks of the model components (i.e., atmosphere and ocean models). In this case, the unaligned land-sea

masks might produce artificial or unrealistic surface heat and momentum fluxes around the coastlines, narrow bays, straits and

seas. The simplest solution is to modify the land-sea masks of the individual model components manually to align them, however, this requires time and is complex (especially when the horizontal grid resolution is high).

Besides, the procedure needs to be repeated each time the model domain (i.e., shift or change in the model domain) or horizontal grid resolution is changed.

Alternative customized interpolation techniques that also include extrapolation help to create more generic and automatized solutions. The RegESM modeling system uses an extrapolation

approach to overcome the mismatched land-sea mask problem for the interaction between atmosphere, ocean and wave components. To perform extrapolation, the driver uses a specialized algorithm to find the mapped and unmapped ocean grid points for every coupling direction (Fig. 4). According to the algorithm, the mapped grid

points have same land-sea mask type

Sophie Valcke 22/11/y 10:25

Supprimé: be configured with

Sophie Valcke 22/11/y 10:26

Supprimé: (coupling type)

Sophie Valcke 22/11/y 10:28

Supprimé: However

Sophie Valcke 22/11/y 10:29

Supprimé: in semi-implicit type coupling (Fig. 3b)

Sophie Valcke 23/11/y 12:02

Supprimé: , the

Sophie Valcke 23/11/y 12:07

Supprimé: to this issue

Sophie Valcke 23/11/y 12:09

Supprimé: .

Sophie Valcke 23/11/y 12:09

Supprimé: H

Sophie Valcke 23/11/y 12:09

Supprimé: the main disadvantage of

Sophie Valcke 23/11/y 12:09

Supprimé: solution is the

Sophie Valcke 23/11/y 12:09

Supprimé: d

Sophie Valcke 23/11/y 12:09

Supprimé: the difficulty to fix the land-sea masks of the . Unlike manual editing of the land-sea masks, land-sea masks,

Sophie Valcke 23/11/y 12:10

Supprimé: same

Sophie Valcke 23/11/y 12:10

Supprimé: .

Sophie Valcke 23/11/y 12:09

Supprimé: whenever

Sophie Valcke 23/11/y 12:11

Supprimé: As a . Unlike manual editing of the land-sea masks, land-sea masks,

Sophie Valcke 23/11/y 12:13

Supprimé: Unlike manual editing of the land-sea masks,

Sophie Valcke 22/11/y 10:50

Supprimé: support

Sophie Valcke 22/11/y 10:50

Supprimé: s

Sophie Valcke 23/11/y 12:13

Supprimé: in the interpolationstage

in both model components (i.e., both are sea or land). On the other hand, the land-sea mask type does not match completely in the case of unmapped grid points. The algorithm first interpolates the field from source to destination grid using grid points just

over the sea and nearest-neighbor type interpolation (from Field_A to Field_B). In this case, if the source field (Field_A) belongs to the ATM component, then the nearest source-to-destination method is used. In other cases, the interpolation is performed using the nearest destination-to-source method. Similarly, the same operation is also performed by using bilinear type interpolation (from Field_A to Field_C). Then, the results of both interpolation (Field_B and Field_C) is compared to find mapped and unmapped grid points and create a new modified mask for the exchange fields (Fig. 4).

After finding mapped and unmapped grid points, the field can be interpolated from the source to the destination grid using a two-step interpolation approach. In the first step, the field is interpolated from source to destination grid using a bilinear interpolation and the original land-sea mask. Then, nearest neighbor type interpolation is used on the destination grid to fill unmapped grid points. One of the main drawbacks of this method is that the result field might include unrealistic values and sharp gradients in the areas of complex land-sea mask structure (i.e., channels, straits). The artifacts around the coastlines can be fixed by applying a light smoothing after interpolation or using more sophisticated extrapolation techniques such as the sea-over-land approach

(Kara et al., 2007; Dominicis et al., 2014), which are not included in the current version of the modeling system. Also, the usage of the mosaic grid along with second-order conservative interpolation method, which gives smoother results when the ratio between horizontal grid resolutions of the source and destination grids are high, can overcome unaligned land-sea mask problem. The next major release of ESMF library (8.0) will include the creep fill strategy (Kara et al., 2007) to fill unmapped grid points.

3 Integration of a Co-processing Component in RegESM Modeling System

The newly designed modeling framework is a combination of the ParaView co-processing plugin – which is called Catalyst (Fabian et al., 2011) – and ESMF library that is specially designed for coupling different ESMs to create more complex regional and global modeling systems. In conventional co-processing enabled simulation systems (single physical model

component such as atmosphere along with co-processing support), the Catalyst is used to integrate ParaView visualization pipeline with the simulation code to support in-situ visualization through the use of application-specific custom adaptor code

(Malakar et al., 2012; Ahrens et al., 2014; O’Leary et al., 2016; Woodring et al., 2016). A visualization pipeline integrates a data flow network in which computation is described as a collection of executable modules that are connected in a directed graph representing how data moves between modules (Moreland, 2013). There are three types of modules in a visualization

pipeline: sources (file readers and synthetic data generators), filters (transforms data), and sinks (file writers and rendering

module that provide images to a user interface). The adaptor code acts as a wrapper layer and transforms information coming

from NUOPC cap to the co-processing component in a compatible format that is defined using ParaView Catalyst and VTK

(Visualization Toolkit) APIs. Moreover, the adaptor code is responsible for defining the underlying computational grid and associating them with the multi-dimensional fields. After defining computational grids and fields, ParaView processes the

received data to perform co-processing to create desired products such as rendered visualizations, added value information (i.e., spatial and temporal averages, derived fields) as well as writing raw data to the disk storage (Fig. 5a).

The implemented novel approach aims to create a more generic and standardized co-processing environment designed explicitly

for Earth system science (Fig. 5b). With this approach, existing ESMs, which are coupled with ESMF library using NUOPC interface, may benefit from the use of an integrated modeling framework to analyze the data

Sophie Valcke 22/11/y 10:55

Supprimé:

Sophie Valcke 22/11/y 10:55

Supprimé:

Sophie Valcke 22/11/y 10:55

Supprimé:

Sophie Valcke 22/11/y 10:55

Supprimé:

Sophie Valcke 22/11/y 11:01

Supprimé: between the source and destination grids

Sophie Valcke 22/11/y 11:02

Supprimé: exchange

Sophie Valcke 22/11/y 11:02

Supprimé: exchange

Sophie Valcke 23/11/y 12:27

Supprimé: result field

Sophie Valcke 23/11/y 12:27

Supprimé: using nearest neighbor type interpolation that is performed in the destination grid (from mapped grid points to unmapped grid points)

Sophie Valcke 23/11/y 12:39

Supprimé: with Co-processing Component

Sophie Valcke 23/11/y 12:38

Supprimé: ,

Sophie Valcke 23/11/y 12:38

Supprimé: ,

Sophie Valcke 22/11/y 18:37

Supprimé: the

Sophie Valcke 23/11/y 12:40

Supprimé: By

Sophie Valcke 23/11/y 12:42

Supprimé: the

Sophie Valcke 23/11/y 12:41

Supprimé: ight

Sophie Valcke 23/11/y 12:41

Supprimé: to

flowing from the multi-component and multi-scale modeling system without extensive code development and restructuring. In this design, the adaptor code interacts with the driver through the use of NUOPC cap and provides an abstraction layer for the co-processing component. As discussed previously, the ESMF framework uses a standardized interface (initialization, run and finalize routines) to plug new model components into existing modeling system such as RegESM in an efficient and optimized way. To that end, the new approach will benefit from the standardization of common tasks in the model components to integrate co-processing component with the existing modeling system. In this case, all information (grids, fields, and metadata) required by ParaView Catalyst is received from the driver, and direct interaction between other model components and the co-processing component is not allowed (Fig. 5b). The implementation logic of the adaptor code is very similar to the conventional approach (Fig. 5a). However, in this case, it uses the standardized interface of the ESMF 5 framework and NUOPC layer to define the computational grid and associated two and three-dimensional fields of model components. The adaptor layer maps the field (i.e., ESMF_Field) and grid (i.e., ESMF_Grid) objects to their VTK equivalents through the use of VTK and co-processing APIs, which are provided by ParaView and co-processing plugin (Catalyst). Along with the usage of the new approach, the interoperability between simulation code and in-situ visualization system are enhanced and standardized. The new design provides an easy-to-develop, extensible and flexible modeling environment for Earth system science.

The development of the adaptor component plays an essential role in the overall design and performance of the integrated modeling environment. The adaptor code mainly includes a set of functions for the initialization (defining computational grids and associated input ports), run and finalize the co-processing environment. Similarly, the ESMF framework also uses the same approach to plug new model components into the modeling system as ESMF components. In ESMF framework, the simulation code is separated into three essential components (initialization, run and finalize) and calling interfaces are triggered by the driver to control the simulation codes (i.e., atmosphere and ocean models). In this case, the initialization phase includes definition and initialization of the exchange variables, reading input (initial and boundary conditions) and configuration files and defining the underlying computational grid (step 1 in Fig. 6). The run phase includes a time stepping loop to run the model component in a defined period and continues until simulation ends (step 4 in Fig. 6). The time interval to exchange data between model and co-processing component can be defined using coupling time step just like the interaction among other model components. According to the ESMF convention, the model and co-processing components are defined as a gridded component while the driver is a coupler component. In each coupling loop, the coupler component prepares exchange fields according to the interaction among components by applying re-gridding (except coupling with co-processing component), performing a unit conversion and common operations over the fields (i.e., rotation of wind field). In the new version of the RegESM modeling system (1.1), the driver is extended to redistribute two and three-dimensional fields from physical model components to allow interaction with the co-processing component. In the initialization phase, the numerical grid of ESMF components is transformed into their VTK equivalents using adaptor code (step 3 in Fig. 6). In this case, ESMF_Grid object is used to create vtkStructuredGrid along with their modified parallel two-dimensional decomposition configuration, which is supported by ESMF/NUOPC grid transfer capability (Fig. 7). According to the design,

Sophie Valcke 23/11/y 12:43
Supprimé: the
Sophie Valcke 23/11/y 12:43
Supprimé: associated with them
Sophie Valcke 23/11/y 12:43
Supprimé: ,

Sophie Valcke 23/11/y 12:44
Supprimé: also ensures
Sophie Valcke 23/11/y 12:44
Supprimé:
Sophie Valcke 23/11/y 12:44
Supprimé:
Sophie Valcke 23/11/y 12:44
Supprimé: integrated

each model component transfers their numerical grid representation to co-processing component at the beginning of the simulation (step 1 in Fig. 6) while assigning independent two-dimensional decomposition ratio to the retrieved grid definitions. The example configuration in Figure 7 demonstrates mapping of 2x3 decomposition ratio (in x and y-direction) of ATM component to 2x2 in COP component. Similarly, the ocean model transfers its numerical grid with 4x4 decomposition ratio to co-processing component with 2x2 (Fig. 7). In this case, ATM and OCN model components do not need to have the same geographical domain.

The only limitation is that the domain of ATM model component must cover the entire OCN model domain for an ATM-OCN coupled system to provide the surface boundary condition for OCN component. The main advantage of the generic implementation of the driver component is to assign different computational resources to the components. The computational resource with accelerator support (GPU) can be independently used by co-processing component to do rendering (i.e., iso-surface extraction, volume rendering, and texture mapping) and processing the high volume of data in an efficient and optimized way.

The initialization phase is also responsible for defining exchange fields that will be transferred among the model components and maps ESMF_Field representations as vtkMultiPieceDataSet objects in co-processing component (step 2-3 in Fig. 6). Due to the modified two-dimensional domain decomposition structure of the numerical grids of the simulation codes, the adaptor code also modifies the definition of ghost regions – a small subset of the global domain that is used to perform numerical operations around edges of the decomposition elements. In this case, the ghost regions (or halo regions in ESMF convention) are updated by using specialized calls, and after that, the simulation data are passed (as vtkMultiPieceDataSet) to the co-processing component. During the simulation, the co-processing component of the modeling system also synchronizes with the simulation code and retrieves updated data (step 5 in Fig. 6) to process and analyze the results (step 6 in Fig. 6). The interaction between driver and the adaptor continues until the simulation ends (step 4, 5 and 6 in Fig. 6) and the driver continues to redistribute the exchange fields using ESMF_FieldRedist calls. The NUOPC cap of model components also supports vertical interpolation of the three-dimensional exchange fields to height (from terrain-following coordinates of RegCM atmosphere model) or depth coordinate (from s-coordinates of ROMS ocean model) before passing information to the co-processing component (COP). Then, finalizing routines of the model and co-processing components are called to stop the model simulations and the data analysis pipeline that destroy the defined data structure/s and free the memory (step 7-8 in Fig. 6).

4 Use Case and Performance Benchmark

To test the capability of the newly designed integrated modeling system, described briefly in the previous section, the three-component (atmosphere, ocean, and co-processing) configuration of RegESM1.1 modeling system is implemented to analyze category 5 Hurricane Katrina. Hurricane Katrina was the costliest natural disaster and has been named one of the five deadliest hurricanes in the history of the United States, and the storm is currently ranked as the third most intense United States land falling tropical cyclone. After establishing [itself](#) in the southern Florida coast as a weak category 1 storm near 22:30 UTC 25 August 2005, it strengthened to a category 5 storm by 12:00 UTC 28 August 2005 as the storm entered the central Gulf of Mexico (GoM). The model simulations are performed [over a 3-day period i.e.](#) 27-30 Aug. 2005, which is the most intense period of the cyclone, to observe the evolution of the Hurricane Katrina and understand the importance of air-sea interaction regarding

Sophie Valcke 23/11/y 14:27

Supprimé: that is

Sophie Valcke 23/11/y 14:27

Supprimé: s

Sophie Valcke 23/11/y 14:28

Supprimé: hed

Sophie Valcke 23/11/y 14:28

Supprimé: between

Sophie Valcke 23/11/y 14:29

Supprimé: for

... [1]

its development and predictability. The next section mainly details the three-component configuration of the modeling system as well as the computing environment, preliminary benchmark results performed with limited computing resource (without GPU support), and analysis of the evolution of Hurricane Katrina.

4.1 Working Environment

The model simulations and performance benchmarks are done on a cluster (SARIYER) provided by the National Center for High-Performance Computing (UHem) in Istanbul, Turkey. The CentOS 7.2 operating system installed in compute nodes are configured with a two Intel Xeon CPU E5-2680 v4 (2.40GHz) processor (total 28 cores) and 128 5 GB RAM. In addition to the compute nodes, the cluster is connected to a high-performance parallel disk system (Lustre) with 349 TB storage capacity. The performance network, which is based on Infiniband FDR (56 Gbps) is designed to give the highest performance for the communication among the servers and the disk system. Due to the lack of GPU accelerators in the entire system, the in-situ visualization integrated performance benchmarks are done with the support of software rendering provided by Mesa library. Mesa is an open-source OpenGL implementation that supports a wide range of graphics hardware each with its back-end called a renderer. Mesa also provides several software-based renderers for use on systems without graphics hardware. In this case, ParaView is installed with Mesa support to render information without using hardware-based accelerators.

4.2 Domain and Model Configurations

The Regional Earth System Model (RegESM 1.1) is configured to couple atmosphere (ATM; RegCM) and ocean (OCN; ROMS) models with newly introduced in-situ visualization component (COP; ParaView Catalyst version 5.4.1) to analyze the evolution of Hurricane Katrina and to assess the overall performance of the modeling system. In this case, two atmospheric model domains were designed for RegCM simulations using one-way nesting approach, as shown in Fig. 8. The outer atmospheric model domain (low-resolution; LR) with a resolution of 27-km is centered at 77.5°W, 25.0°N and covers almost entire the United States, the western part of Atlantic Ocean and north-eastern part of Pacific Ocean for better representation of the large-scale atmospheric circulation systems. The outer domain is enlarged as much as possible to minimize the effect of the lateral boundaries of the atmospheric model in the simulation results of the inner model domain (high-resolution; HR). The horizontal grid spacing of inner domain is 3-km and covers the entire GoM and the western Atlantic Ocean to provide high resolution atmospheric forcing for coupled atmosphere-ocean model simulations and perform cloud-resolving simulations. Unlike the outer domain, the model for the inner domain is configured to use the non-hydrostatic dynamical core (available in RegCM 4.6) to allow better representation local scale vertical acceleration and essential pressure features. The lateral boundary condition for the outer domain is obtained from European Centre for Medium-Range Weather Forecasts (ECMWF) latest global atmospheric reanalysis (ERA-Interim project; Dee et. al., 2011), which is available at 6-h intervals at a resolution of 0.75°x0.75° in the horizontal and 37 pressure levels in the vertical. On the other hand, the lateral boundary condition of the HR domain is specified by the results of the LR domain. Massachusetts Institute of Technology-Emanuel convective parameterization scheme (MIT-EMAN; Emanuel, 1991; Emanuel and Zivkovic-Rothman, 1999) for the cumulus representation along with sub-grid explicit moisture (SUBEX; Pal et al., 2000) scheme for large-scale precipitations are used for LR outer domain.

As it can be seen in Fig. 8, the ROMS ocean model is configured to cover entire the GoM to allow better tracking of the Hurricane Katrina. In this case, the used ocean model configuration is very similar to the configuration used by Physical

Sophie Valcke 23/11/y 14:30

Supprimé: includes

Sophie Valcke 23/11/y 14:30

Supprimé: ed

Sophie Valcke 23/11/y 14:30

Supprimé: information of

Sophie Valcke 23/11/y 14:29

Supprimé:

Sophie Valcke 23/11/y 14:29

Supprimé: s

Sophie Valcke 23/11/y 14:31

Supprimé: used

Sophie Valcke 23/11/y 14:31

Supprimé: the

Sophie Valcke 23/11/y 14:31

Supprimé: that are done in

Sophie Valcke 23/11/y 14:34

Supprimé: Concerning cumulus convection,

Sophie Valcke 23/11/y 14:35

Supprimé: is used in outer model simulations.

Sophie Valcke 23/11/y 14:35

Supprimé: A

Sophie Valcke 23/11/y 14:35

Supprimé: selected cumulus convection parameterization,

Sophie Valcke 23/11/y 14:36

Supprimé: is used to represent

Oceanography Numerical Group (PONG), Texas A&M University (TAMU), in which the original model configuration can be accessed from their THREDDS server. The ocean model has a spatial resolution of $1/36^\circ$, which corresponds to a non uniform resolution of around 3 km (655 x 489 grid points) with highest grid resolution in the northern part of the domain. The model has 60 vertical sigma layers ($_s = 10.0, _b = 2.0$) to provide detailed representation of the main circulation patterns of the region and vertical tracer gradients. The bottom topography data of the GoM is constructed using 5 the ETOPO1 dataset (Amante and Eakins, 2009), and minimum depth (hc) is set to 400 m. The bathymetry js also modified so that the ratio of depth of any two adjacent columns does not exceed 0.25 to enhance the stability of the model and ensure hydrostatic consistency that prevents pressure gradient error. The Mellor-Yamada level 2.5 turbulent closure (MY; Mellor and Yamada, 1982) is used for vertical mixing, while rotated tensors of the harmonic formulation are used for horizontal mixing. The lateral boundary conditions for ROMS ocean model are provided by Naval Oceanographic Office Global Navy Coastal Ocean Model (NCOM) during 27-30 August 2005. The model coupling time step between atmosphere and ocean model component is set to 1 hour but 6 minutes coupling time step is used to provide one-way interaction with co-processing component to study Hurricane Katrina in a very high temporal resolution. In the coupled model simulations, the ocean model provides SST data to the atmospheric model in the region where their numerical grids overlap. In the rest of the domain, the atmospheric model uses SST data provided by ERA-Interim dataset (prescribed SST). The results of the performance benchmark also include additional tests with smaller coupling time step such as 3 minutes for the interaction with the co-processing component. In this case, the model simulations for the analysis of Hurricane Katrina runs over three days, but only one day of simulation length is chosen in the performance benchmarks to reduce the compute time.

4.3 Performance Benchmark

A set of simulations are performed with different model configurations to assess the overall performance of the coupled modeling system by focusing on the overhead of the newly introduced co-processing component (Table 1). The performance benchmarks include analysis of the extra overhead provided by the co-processing component, coupling interval between physical models and co-processing component under different rendering load such as various visualization pipelines (Table 1). Two different atmospheric model configurations (a low-resolution, LR and high-resolution HR) are defined to scale up to a large number of processors. The LR model domain includes around 900.000 grid points in the atmospheric model while the HR domain contains 25 million grid points. In both cases, the ocean model configuration is the same, and it has around 19 million grid points. Besides the use of a non-hydrostatic dynamical core in the atmospheric model component in the HR case, the rest of the model configuration is preserved. To isolate the overhead of the driver from the overhead of the co-processing component, first individual model components (ATM and OCN) are run in standalone mode and then, the best-scaled model configurations regarding two-dimensional decomposition configuration are used in the coupled model simulations (CPL and COP). Due to the current limitation in the integration of the co-processing component, the coupled model only supports sequential type execution (see Section 2.5 for more information) when the co-processing component is activated, but this limitation will be removed in the future version of the modeling system (RegESM2.0). As mentioned in the previous section, the length of the simulations is kept relatively short (1 day) in the benchmark analysis to perform many simulations with different model configurations (i.e., coupling interval, visualization pipelines and domain decomposition parameters). The benchmark results of standalone model components (ATM and OCN) can be seen in Fig. 9. In this case, two different atmospheric model configurations are considered to see the effect of the domain size and non-

Sophie Valcke 23/11/y 14:39
Supprimé: data are
 Sophie Valcke 23/11/y 14:39
Supprimé: s
 Sophie Valcke 23/11/y 14:39
Supprimé: grid
 Sophie Valcke 23/11/y 14:39
Supprimé: creation

Sophie Valcke 23/11/y 14:41
Supprimé: change
 Sophie Valcke 23/11/y 14:41
Supprimé: the
 Sophie Valcke 23/11/y 14:41
Supprimé: of
 Sophie Valcke 23/11/y 14:42
Supprimé: in
 Sophie Valcke 23/11/y 14:42
Supprimé: (non-hydrostatic)
 Sophie Valcke 23/11/y 14:42
Supprimé: s are

hydrostatic dynamical core in the benchmark results (LR and HR; Fig. 8). The results show that the model scales pretty well and it is clear that the HR case shows better scaling results than LR configuration of the atmospheric component (ATM) as expected. It is also shown that around 588

processors, which is the highest available compute resource, the communication among the processors dominate the benchmark

results and even HR case does not gain further performance (Fig. 9a). Similar to the atmospheric model component, the ocean

model (OCN) is also tested to find the best two-dimensional domain decomposition configuration (tiles in x and y-direction).

As it can be seen from the Fig. 9b, the selection of the tile configuration affects the overall performance of the ocean model.

In general, model scales better if the tile in the x-direction is bigger than the tile in the y-direction, but this is more evident in the

small number of processors. The tile effect is mainly due to the memory management of Fortran programming language (column-major order) as well as the total number of active grid points (not masked as land) placed in each tile. Tile options must be selected carefully while considering the dimension of the model domain in each direction.

In some tile configuration, it is not possible to run the model due to the underlying numerical solver and the required

minimum ghost points. To summarize, the ocean model scales well until 588 cores with the best tile configurations indicated in Fig. 9b.

The performance of the two-component modeling system (CPL) can be investigated using the benchmark results of the

standalone atmosphere and ocean models. In this case, the best two-dimensional decomposition parameters of the standalone ocean model simulations are used in the coupled model simulations (Fig. 9b). The comparison of the standalone and coupled

model simulations shows that the driver component introduces additional 5-10% (average is 5% for LR and 6% for HR cases)

overhead in the total execution time, which slightly increases along with the used total number of processors. This overhead can be considered acceptable

when increased number of MPI communication between the components are considered (Fig. 9 and 10a). The extra overhead

is mainly due to the interpolation (sparse matrix multiply performed by ESMF) and extrapolation along the coastlines to match land-sea masks of the atmosphere and ocean models and fill the unmapped grid points to exchange data (Fig. 4).

To investigate the overhead introduced by the newly designed co-processing component, the three-component modeling

system (COP) is tested with three different visualization pipelines (P1, P2, and P3; Table 1) using two different atmospheric

model configurations (LR and HR) and coupling interval (3 and 6 minutes with co-processing). In this case, the measured

total execution time during the COP benchmark results also includes vertical interpolation (performed in ESMF cap) to map data from sigma coordinates to height (or depth) coordinates for both physical model components (ATM and OCN). As shown

in Fig. 10b-d, the co-processing components require 10-40% extra execution time for both LR and HR cases depending on

used visualization pipeline when it is compared with CPL simulations. The results also reveal that the fastest visualization

pipeline is P3 and the slowest one is P1 for the HR case (Fig. 10b and d). Table 1 also includes the execution time of the single visualization pipeline (measured by using MPI_Wtime call) isolated from the rest of the tasks.

In this case, each rendering task

gets 2-4 seconds for P1 and P2 cases and 7-15 seconds for the P3 case in LR atmospheric model configuration. For HR case, P1 and P2 take around 17-80 seconds, and the P3 case is rendered in around 8-10 seconds. These results show that the time spent

in the co-processing component (sending data to ParaView, Catalyst, and rendering to create the output) fluctuates too much and that this component does not present a

predictable and stable behavior. It might be due to the particular configuration of the ParaView, which is configured

Sophie Valcke 23/11/y 14:52

Supprimé: in the results

Sophie Valcke 23/11/y 14:53

Supprimé: On the

... [2]

Sophie Valcke 23/11/y 14:53

Supprimé: t

Sophie Valcke 23/11/y 14:53

Supprimé: used

Sophie Valcke 23/11/y 14:59

Supprimé: ,

Sophie Valcke 23/11/y 14:59

Supprimé: which is

Sophie Valcke 23/11/y 15:12

Supprimé: do not show

Sophie Valcke 23/11/y 15:12

Supprimé: ,

to use software-based rendering to process data in CPUs and load in the used high-performance computing system (UHeM) even if the benchmark tests are repeated multiple times.

In addition to the testing modeling system with various data processing load, a benchmark with increased coupling time step is also performed (see P23M in Fig. 10c). In this case, the coupling time step between physical model components and the co-processing component is increased (from 6 minutes to 3 minutes) to produce output in doubled frame rate, but coupling interval between physical model components (ATM and OCN) are kept same (1 hour). The benchmark results show that increased coupling time step also rises overhead due to the co-processing from 45% to 60% for HR case and pipeline P2 when it is compared with the results of two-component simulations (CPL; Fig. 10a). It is also shown that the execution time of co-processing enabled coupled simulations increase but the difference between P2 and P23M cases are reduced from 66% to 37% when the number of processors increased from 140 to 588.

In addition to the analysis of timing profiles of modeling system under different rendering load, the amount of data exchanged and used in the in-situ visualization case can be compared with the amount of data that would be required for offline visualization at the same temporal frequency to reveal the added value of the newly introduced co-processing component. For this purpose, the amount of data exchanged with co-processing component is given in Table 1 for three different visualization pipelines (P1, P2, and P3). In co-processing mode, the data retrieved from model component **memory** (single time step) **by the driver is passed to ParaView Catalyst** for rendering. **In addition to processing data concurrently with the simulation on co-processing component, the offline visualization (post-processing) consists of computations** done after the model is run and requires to store numerical results in a disk environment. For example, 3-days long simulation with 6-minutes coupling interval produces around 160 GB data (720 time-step) just for a single variable from high-resolution atmosphere component (P1 visualization pipeline) in case of using offline visualization. **With co-processing, the same analysis can be done by applying the same visualization pipeline (P1), which requires to process only 224 MB data stored in the memory, in each coupling interval. Moreover, storing results of three-day long, high-resolution simulation of RegCM atmosphere model (in netCDF format) for offline visualization requires around 1.5 TB data in case of using 6-minutes interval in the default configuration (7 x 3d fields and 28 x 2d fields). It is evident that the usage of co-processing component reduces the amount of data stored in the disk and allows more efficient data analysis pipeline.**

Besides the minor fluctuations in the benchmark results, the modeling system with co-processing component scales pretty well to the higher number of processors (or cores) without any significant performance pitfalls in the current configuration.

On the other hand, the usage of accelerator enabled ParaView configuration (i.e., using NVIDIA EGL library) and ParaView plugins with accelerator support such as NVIDIA IndeX volume rendering plugin and new VTK-m filters to process data on GPU will improve the benchmark result. The NVIDIA IndeX for ParaView Plugin enables large-scale and high-quality volume data visualization capabilities of the NVIDIA IndeX library inside the ParaView and might help to reduce time to process highresolution spatial data (HR case). In addition to NVIDIA IndeX plugin, VTK-m is a toolkit of scientific visualization algorithms for emerging processor architectures such as GPUs (Moreland, 2016). The model configurations used in the simulations also write simulation results to the disk in netCDF format. In case of disabling of writing data to disk or configure the models to write data with large time intervals (i.e., monthly), the simulations with active co-processing component will run much faster and make the analysis of the model results in real time efficiently especially in live mode (see Section 5.1).

- Sophie Valcke 23/11/y 17:47
Supprimé: through the use [3]
- Sophie Valcke 23/11/y 17:47
Supprimé: read from memory and
- Sophie Valcke 23/11/y 17:47
Supprimé: the
- Sophie Valcke 23/11/y 17:47
Supprimé: ,
- Sophie Valcke 23/11/y 17:56
Supprimé: Besides,
- Sophie Valcke 23/11/y 17:49
Supprimé: of the
- Sophie Valcke 23/11/y 17:50
Supprimé: that are
- Sophie Valcke 23/11/y 15:16
Supprimé: In the case of
- Sophie Valcke 23/11/y 15:17
Supprimé: using

5 Demonstration Application

The newly designed modeling system can analyze numerical simulation results in both in-situ (or live) and co-processing modes.

In this case, a Python script, that defines the visualization pipeline, mainly controls the selection of the operating mode and is generated using ParaView, co-processing plugin. The user could also activate live visualization mode, just by changing a single line of code (need to set `coprocessor.EnableLiveVisualization` as `True`) in Python script. This section aims to give more detailed information about two different approaches by evaluating numerical simulation of Hurricane Katrina in both models to reveal the designed modeling system capability and its limitations.

5.1 Live Visualization Mode

While the live visualization designed to examine the simulation state at a specific point in time, the temporal filters such as ParticlePath, ParticleTracer, TemporalStatistics that are designed to process data using multiple time steps cannot be used in

this mode. However, live visualization mode allows connecting to the running simulation anytime through the ParaView GUI

in order to make detailed analysis by modifying existing visualization pipelines defined by a Python script. In this case, the

numerical simulation can be paused while the visualization pipeline is modified and will continue to run with the revised one. It is evident that the live visualization capability gives full control to the user to make further investigation about the simulation results and facilitate better insight into the underlying physical process and its evolution in time.

The current version of the co-processing enabled modeling system can process data of multiple model components by

using multi-channel input port feature of ParaView Catalyst. In this case, each model has two input channels based on the

rank of exchange fields. For example, atmospheric model component has `atm_input2d` and `atm_input3d` input channels to make processing available both two and three-dimensional exchange fields. The underlying adaptor code resides between the

NUOPC cap of co-processing component and ParaView Catalyst and provides two grid definitions (2d and 3d) for each model

components for further analysis. In this design, the ParaView Co-processing Plugin is used to generate Python co-processing scripts, and user needs to map data sources to input channels by using predefined names such as `atm_input2d` and `ocn_input3d`.

Then, adaptor provides the required data to co-processing component through each channel to perform rendering and data analysis in real time. The fields that are used in the co-processing component are defined by generic ASCII formatted driver configuration file (`exfield.tbl`), which is also used to define exchange fields among other model components such as atmosphere

and ocean models. Fig. 11 shows a screenshot of live visualization of three-dimensional relative humidity field provided by the low-resolution atmospheric model component, underlying topography information, and vorticity of ocean surface that is provided by ocean model component.

5.2 Co-processing Mode

In addition to live visualization mode that is described briefly in the previous section, ParaView Catalyst also allows to render

and store data using predefined co-processing pipeline (in Python) for further analysis. Co-processing mode can be used for

three purposes: (1) the simulation output can be directed to the co-processing component to render data in batch mode and write image files to the disk, (2) added value information (i.e., vorticity from wind components, eddy kinetic energy from

Sophie Valcke 23/11/y 16:05

Supprimé: available

Sophie Valcke 23/11/y 16:06

Supprimé: ,

ocean current) can be calculated and stored in a disk for further analysis and (3) storing simulation output in a higher temporal resolution to process it later (post-processing) or create a representative dataset that can be used to design visualization pipeline for co-processing or live visualization modes. In this case, the newly designed modeling system can apply multiple visualization and data processing pipelines to the simulation results at each coupling time step to make a different set of analysis over the results of same numerical simulation for more efficient data analysis. The modeling system also facilitates multiple input ports to process data flowing from multiple ESM components. In this design, input ports are defined automatically by the co-processing component based on activated model components (ATM, OCN, etc.) and each model component has two ports to handle two and three-dimensional grids (and fields) separately such as atm_input2d, atm_input3d, ocn_input2d and ocn_input3d. To test the capability of the co-processing component, the evolution of Hurricane Katrina is investigated by using two different configurations of the coupled model (COP_LR and COP_HR) that are also used to analyze the overall computational performance of the modeling system (see Section 4.3). In this case, both model configurations use the same configuration of OCN model component, but the different horizontal resolution of the ATM model is considered (27 km for LR and 3 km for HR cases). Figure 12 shows 3-hourly snapshots of the model simulated clouds that are generated by processing three-dimensional relative humidity field calculated by the low-resolution version of the coupled model (COP_LR) using NVIDIA IndeX volume rendering plugin as well as streamlines of Hurricane Katrina, which is calculated using three-dimensional wind field. The visualization pipeline also includes sea surface height and surface current from the ocean model component to make an integrated analysis of the model results. Figure 12a-b shows the streamlines that are produced by extracting the hurricane using ParaView Threshold filter. In this case, the extracted region is used as a seed to calculate backward and forward streamlines. In Figure 12c-e, sea surface height, sea surface current and surface wind vectors (10-meters) are shown together to give insight about the interaction of ocean-related variables with the atmospheric wind. Lastly, the hurricane reaches to the land and starts to disappear due to increased surface roughness and lack of energy source (Fig. 12f). While the low-resolution of atmosphere model configuration is used, the information produced by the new modeling system enabled to investigate the evolution of the hurricane in a very high temporal resolution, which was impossible before. A day-long animation that is also used to create Figure 12 can be found as a supplemental video (Turuncoglu, 2018a). In addition low-resolution model results revealing the evolution of the hurricane in a very high temporal resolution, low and high-resolution model results are also compared to see the added value of the increased horizontal resolution of the atmospheric model component regarding representation of the hurricane and its structure. To that end, a set of visualization pipelines are designed to investigate the vertical updraft in the hurricane, simulated track, precipitation pattern, and ocean state. In this case, two time snapshots are considered: (1) 28 August 2005 0000 UTC, at the early stage of the hurricane in Category 5 and (2) 29 August 2005 0000 UTC just before Katrina makes its third and final landfall near Louisiana–Mississippi border, where the surface wind is powerful, and surface currents had a strong onshore component (McTaggart-Cowan et al., 2007a, b). In the analysis of vertical structure, the hurricane is isolated based on the criteria of surface wind speed that exceeds 20 ms⁻¹ and the seed (basically set of points defined as vtkPoints) for ParaView StreamTracerWithCustomSource filter are defined dynamically using ProgrammableFilter as a circular plane with a radius of 1.2° and points distributed with 0.2° interval in both direction (x and y) around the center of mass of the isolated

Sophie Valcke 23/11/y 16:07

Supprimé: s

Sophie Valcke 23/11/y 16:07

Supprimé: have

Sophie Valcke 23/11/y 16:07

Supprimé: s

Sophie Valcke 23/11/y 16:11

Supprimé: to the analysis of

Sophie Valcke 23/11/y 16:11

Supprimé: to

Sophie Valcke 23/11/y 16:11

Supprimé: when it is

Sophie Valcke 23/11/y 16:11

Supprimé: when it is

region. Then, forward and backward streamlines of vorticity are computed separately to see inflow at low and mid levels and outflow at upper levels for both low (COP_LR; Fig. 13a, b, d and e) and high-resolution (COP_HR; Fig. 14a, b, d and e) cases. The analysis of simulations reveal that the vertical air movement shows higher spatial variability in high-resolution simulation (COP_HR) case

even if the overall structure of the hurricane is similar in both cases. As expected, the strongest winds occur in a region forming a ring around the eyewall of the hurricane, which is where the lowest surface pressure occurs. In addition, the analysis

of cloud liquid water content also shows that low and mid-levels of the hurricane have higher water content in a decreasing trend with height and spatial distribution of precipitation is better represented in high resolution case (Fig. 14a-b and d-e),

which is consistent with the previous modelling study of Trenberth et al. (2007). It is also seen that the realistic principal and secondary precipitation bands around the eye of the hurricane are more apparent and well structured in the high-resolution

simulation while the low-resolution case does not show those small-scale features (Fig. 13a-b and d-e). On the ocean side, the loop current, which is a warm ocean current that flows northward between Cuba and the Yucatan Peninsula and moves north into the Gulf of Mexico, loops east and south before exiting to the east through the Florida Straits and joining the Gulf Stream and is well defined by the ocean model component in both cases (Fig. 13c and f;

Fig. 14c and f). The track of the hurricane is also compared with the HURDAT2 second-generation North Atlantic (NATL) hurricane database, which is the longest and most complete record of tropical cyclone (TC) activity in any of the world's oceans

(Landsea and Franklin, 2013). In this case, the eye of the hurricane is extracted as a region with surface pressure anomaly greater than 15 millibar (shown as a circular region near the best track). As it can be seen from the figures, Katrina moves over in the

central Gulf, which is mainly associated with the loop current and persistent warm and cold eddies, and intensifies as it passes

over the region due to the high ocean heat content in both simulation (Fig. 13c and f and Fig. 14c and f). The comparison of the low and high-resolution simulations also indicate that the diameter of hurricane-force winds at peak intensity is bigger in high-resolution simulation case at 29 August 2005 0000 UTC (Fig. 13f and Fig. 14f). An animation that shows the comparison of low and high-resolution model results can be found as a supplemental video (Turuncoglu, 2018b).

While the main aim of this paper is to give design details of the new in-situ visualization integrated modeling system and show its capability, the performance of the coupled modeling system to represent one of the most destructive hurricanes is very satisfactory especially for high-resolution case (COP_HR). Nonetheless, the individual components (atmosphere and ocean) of

the modeling system can be tuned to have better agreement with the available observations and previous studies. Specifically for the analysis of the hurricane, a better storm tracking algorithm needs to be implemented using ParaView Programmable

Filter by porting existing legacy Fortran codes for more accurate storm tracking in both live and co-processing mode.

6 Discussion of The Concepts Associated with Interoperability, Portability, and Reproducibility

In the current design of the RegESM modeling system, the NUOPC cap of the co-processing component is designed to work with regional modeling applications that have specific horizontal grid (or mesh) types such as rectilinear and curvilinear grids.

The newly introduced co-processing interface (NUOPC cap and adaptor code) now needs to be generalized to be compatible

with other regional and global modeling systems coupled through ESMF and NUOPC layer. Specifically, the following issues

need to be addressed to achieve better interoperability with existing modeling systems and model components:

(1) redesigning the NUOPC cap of the co-processing component to support various global and regional mesh

Sophie Valcke 23/11/y 16:12

Supprimé: it is

Sophie Valcke 23/11/y 16:12

Supprimé: ed

Sophie Valcke 23/11/y 16:12

Supprimé: as

Sophie Valcke 23/11/y 16:18

Supprimé: In addition to the analysis ... [4]

Sophie Valcke 23/11/y 16:19

Supprimé: ,

Sophie Valcke 23/11/y 16:20

Supprimé: that

Sophie Valcke 23/11/y 16:20

Supprimé: is

Sophie Valcke 23/11/y 16:21

Supprimé: d

Sophie Valcke 23/11/y 16:21

Supprimé: d

Sophie Valcke 23/11/y 16:56

Supprimé: herefore, t

Sophie Valcke 23/11/y 16:56

Supprimé: with

Sophie Valcke 23/11/y 16:57

Supprimé: newly introduced

types such as Cubed-Sphere and unstructured Voronoi meshes, (2) extending the adaptor code to represent mesh and exchange fields provided by NUOPC cap using VTK and ParaView Catalyst APIs, (3) adding support to co-processing interface for models with nesting capability, and (4) adding support to have common horizontal grid definitions in the co-processing component and in the other components to make integrated analysis of data (i.e., calculating air-sea temperature difference and correlation) produced by the various model components. Moreover, the co-processing interface can be tightly integrated with the NUOPC layer to provide a simplified API for designing new in-situ visualization integrated modeling systems in an efficient and standardized way. Besides the configuration used in this study, the RegESM modeling system is also tested with different model configurations such as coupling RegCM, MITgcm, and co-processing component to investigate air-sea interaction in the Black Sea basin. Initial results show that the co-processing component can also successfully process data flowing from different model configurations supported by RegESM.

When the diverse nature of high-performance computing systems, their hardware infrastructure (i.e., performance networks and storage systems) and software stacks (i.e., operating systems, compilers, libraries for inter-node communication and their different versions) are considered, realizing fully portable modeling system becoming increasingly crucial for the scientific community. In this case, the detailed examination of possible configurations of the modeling system and exiting computing environments can help to improve the flexibility and portability of the modeling system. Specifically for RegESM modeling system, the use case application and benchmark simulations revealed that the single executable approach (combining all model components into one program) used in the design of the modeling system can cause a portability problem when visualization and simulation run on concurrent resources. In the case of a homogeneous computing environment (all nodes with or without GPU support), the in-situ enabled modeling system runs without any particular problem because each MPI (Message Passing Interface) process has access to the same software and hardware resources. In contrast, some computing systems may have heterogeneous underlying hardware and software stack (e.g., mixed servers with and without GPU support). As a result, the simulation with in-situ visualization would fail due to missing shared software libraries in underlying GPU. In this case, two approaches can be used to overcome the problem: (1) installation of required libraries on the entire system even on servers that do not have GPU support, and (2) restructuring modeling system to support two executables, one for the co-processing component and one for the physical model component. The second approach is considered a more generic and flexible solution and enhances the portability of the modeling system. It also allows implementing a loosely coupled in-situ visualization system and enables the use of specialized hardware (GPU and more memory) for rendering (Rivi et al., 2012). The main drawback of the loosely coupled in-situ visualization approach is that it requires transferring data over the network. As a result, the network performance can be a bottleneck for the modeling system, especially for high-resolution multi-component modeling applications. When the complexity of regional and global ESMs are considered, developing fully reproducible, and portable modeling system is a challenging task and requires significant human interaction to keep track of detailed metadata and provenance information about the model, simulation and computing environment (in both software and hardware levels). The use of scientific workflows in Earth System science has demonstrated advantages in terms of metadata, provenance, error handling, and reproducibility in an automatized and standardized way (Turuncoglu et al., 2011, 2012; Turuncoglu, 2012). Additionally, the rapid development in the software container technology can help to design flexible and portable computing environments. Hence, the Docker container was implemented to examine the feasibility of using the container approach for our newly developed in-situ visualization integrated modeling system. A

- Sophie Valcke 23/11/y 16:57
Supprimé: ... [5]
- Sophie Valcke 23/11/y 16:58
Supprimé: addressing ... [6]
- Sophie Valcke 23/11/y 16:59
Supprimé: The i
- Sophie Valcke 23/11/y 16:59
Supprimé: rather ... [7]
- Sophie Valcke 23/11/y 16:59
Supprimé: developed
- Sophie Valcke 23/11/y 17:00
Supprimé: es
- Sophie Valcke 23/11/y 17:00
Supprimé: es
- Sophie Valcke 23/11/y 17:00
Supprimé: have
- Sophie Valcke 23/11/y 17:00
Supprimé:
- Sophie Valcke 23/11/y 17:00
Supprimé: to the homogeneous comp... [8]
- Sophie Valcke 23/11/y 17:00
Supprimé: the
- Sophie Valcke 23/11/y 17:01
Supprimé: does not have a uniform view of
- Sophie Valcke 23/11/y 17:01
Supprimé: i.e.,
- Sophie Valcke 23/11/y 17:01
Supprimé: in a heterogeneous computing environment
- Sophie Valcke 23/11/y 17:01
Supprimé: support
- Sophie Valcke 23/11/y 17:02
Supprimé: specific to
- Sophie Valcke 23/11/y 17:02
Supprimé: in all used nodes or servers
- Sophie Valcke 23/11/y 17:02
Supprimé: to
- Sophie Valcke 23/11/y 17:03
Supprimé: s
- Sophie Valcke 23/11/y 17:03
Supprimé: (i.e., RegESM)
- Sophie Valcke 23/11/y 17:04
Supprimé: e
- Sophie Valcke 23/11/y 17:04
Supprimé: s
- Sophie Valcke 23/11/y 17:04
Supprimé: for earth system modeling
- Sophie Valcke 23/11/y 17:04
Supprimé: has been
- Sophie Valcke 23/11/y 17:05
Supprimé: (RegESM).

container is a standard unit of software that helps to create a software package including all its dependencies, which can then be ported from one computing environment to another without worrying about the underlying hardware infrastructure and software stack. It also enhances the numerical reproducibility of simulations by creating a standardized computing environment isolated from any dependencies. In this study, Docker is selected as a container environment because it is widely adopted across the software industry and has a very active user community. Despite the flexibility and easy to use nature of Docker containers, using specialized hardware such as NVIDIA GPUs, which require kernel modules and user-level libraries to operate, is not supported natively. Therefore, Docker container cannot access the underlying GPU resource to perform hardware level rendering for in-situ visualization. To enable portable GPU-based containers, NVIDIA developed a special container that mounts the user mode components of the driver and the GPUs into the container at launch. As a part of this study, the newly developed RegESM modeling system was tested with both Docker (software rendering through the use of Mesa library) and NVIDIA Docker (hardware based rendering). The initial results show that RegESM can take advantage of the container approach to create portable and reproducible modeling system in both in-situ and co-processing modes without considerable performance loss (5-10%). The added value of using NVIDIA Docker is that it enables to utilize the underlying GPU resource to perform rendering (i.e., representation of clouds using direct volume rendering method). More information about creating a Docker container for in-situ visualization enabled modeling system can be found in the dedicated GitHub repository (see code availability section).

7 Summary and Conclusions

In this study, the newly developed in-situ visualization integrated modeling system (RegESM1.1) is used to demonstrate the feasibility and added value of the integrated modeling environment to analyze the high volume of data coming from a multi-component ESM, which was not possible before. In this case, ParaView, Catalyst is used as a co-processing component to process and render data. The results of the selected use case (evolution of Hurricane Katrina) show that the co-processing component provides an easy-to-use and generic modeling and data analysis environment, which is independent of the underlying physical model components used. Moreover, it promotes the usage of co-processing capability with the existing ESMs coupled using ESMF framework and NUOPC layer, without significant code restructuring and development and helps to increase the interoperability between ESMs and ParaView, co-processing plugin (Catalyst). In the current implementation, the prototype version of the adaptor code acts as an abstraction layer to simplify and standardize the regular tasks to integrate the simulation code with in-situ visualization and analysis environment. The driver is also responsible for redistributing the data to the co-processing component while preserving its numerical grid along with the support of vertical interpolation. Coupling of the co-processing component with the generic driver facilitates the definition of custom data processing pipelines (defined by Python scripts) and allows analysis of data originated from different components (i.e., atmosphere and ocean models) of the RegESM modeling system in a very high temporal resolution. In this way, RegESM modeling system can be used to study various physical processes (i.e., extreme precipitation events, air-sea interaction, convection, and turbulence) that could not be analyzed with the conventional post-processing approaches. While the results of the in-situ visualization integrated modeling system are encouraging, the co-processing component will be extended to support different regional and global computational grid representations supported by ESMF library such as

Sophie Valcke 23/11/y 17:06

Supprimé: for software and

Sophie Valcke 23/11/y 17:07

Supprimé: . By this way, the applicat ... [9]

Sophie Valcke 23/11/y 17:07

Supprimé: results

Sophie Valcke 23/11/y 17:07

Supprimé: the

Sophie Valcke 23/11/y 17:08

Supprimé: the

Sophie Valcke 23/11/y 17:08

Supprimé: the

Sophie Valcke 23/11/y 17:08

Supprimé: In this case

Sophie Valcke 23/11/y 17:12

Supprimé: to evaluate the potential of using container technology to improve porta ... [10]

Sophie Valcke 23/11/y 17:12

Supprimé: the developed modeling system

Sophie Valcke 23/11/y 17:13

Supprimé: the

Sophie Valcke 23/11/y 17:16

Supprimé: state-of-art

Sophie Valcke 23/11/y 17:17

Supprimé: in an integrated way

Sophie Valcke 23/11/y 17:17

Supprimé: ,

Sophie Valcke 23/11/y 17:18

Supprimé:

Sophie Valcke 23/11/y 17:18

Supprimé:

Sophie Valcke 23/11/y 17:18

Supprimé: ,

Sophie Valcke 23/11/y 17:18

Supprimé: which is

Sophie Valcke 23/11/y 17:19

Supprimé: ,

Sophie Valcke 23/11/y 17:19

Supprimé: The c

Sophie Valcke 23/11/y 17:20

Supprimé: o

Sophie Valcke 23/11/y 17:20

Supprimé: e

Sophie Valcke 23/11/y 17:20

Supprimé: easily

unstructured meshes for having a generic adaptor for various model applications. Additionally, we are currently exploring: (1) the way to optimize the transfer of grid features and mapping of exchange fields to enhance the overall performance of the modeling environment in terms of memory usage and computational efficiency, especially for very high-resolution applications (< 3 km), (2) possibility of automatic detection of accelerators (GPUs) through the use of driver component and assigning available GPU resources automatically to the co-processing component for rendering, (3) improving modeling system and co-processing component to allow nested applications (both atmosphere and ocean), (4) developing more application of the integrated modeling environment (possibly with other ocean and atmosphere components such as WRF and MITgcm) to analyze different physical processes such as air-sea interactions in upwelling regions under extreme atmospheric forcing conditions.

Code availability.

The RegESM modeling system is open source and available under the MIT License, making it suitable for community usage. The license allows modification, distribution, private and commercial uses. The source code for all versions of RegESM driver including 1.1 is distributed through the public code repository hosted by GitHub (<https://github.com/uturuncoglu/RegESM>). The user guide and detailed information about the modeling system are also distributed along with the source code in the same code repository. The RegESM source code includes required code patches for the individual model components to be used as a component in the modeling system. On the other hand, the source code of individual model components such as the ocean, wave, and river routing components and co-processing tool (ParaView, Catalyst) are distributed mainly by their home institutes that might apply different licensing types. The reader who wants to get more information about the individual model components and their license type can refer to their websites. The release version 1.1 is permanently archived on Zenodo and accessible under the digital object identifier doi:10.5281/zenodo.1307212. The demo configuration of the modeling system that is used in NVIDIA GPU Technology Conference (GTC) 2018 is also permanently archived on Zenodo and accessible under the digital object identifier doi:10.5281/zenodo.1474753. The repository also includes detailed information about the installation of the individual components of the modeling system, third-party libraries, and commands to create Docker container.

Competing interests.

The author declares that he has no conflict of interest.

Acknowledgements.

This study has been supported by a research grant (116Y136) provided by The Scientific and Technological Research Council of Turkey (TUBITAK). The computing resources used in this work were provided by the National Center for High Performance Computing of Turkey (UHeM) under grant number 5003082013 and 5004782017. The Quadro K5200 used for the development of the prototype version of the modeling system was donated by the NVIDIA Corporation as part of the Hardware Donation Program. The author extends his grateful thank to Rocky Dunlap and Robert Oehmke from NOAA/ESRL, and CIRES, Boulder, Colorado, Gerhard Theurich from Science Applications International Corporation, McLean, Virginia and Andy C. Bauer from Kitware Inc., USA and Mahendra Roopa from NVIDIA for their very useful suggestions and comments.

Sophie Valcke 23/11/y 17:21

Supprimé: transfer

Sophie Valcke 23/11/y 17:23

Supprimé: to

Sophie Valcke 23/11/y 17:23

Supprimé: using

Sophie Valcke 23/11/y 17:23

Supprimé: the

Sophie Valcke 23/11/y 17:24

Supprimé: them

Sophie Valcke 23/11/y 17:24

Supprimé: used in the modeling system

Sophie Valcke 23/11/y 17:24

Supprimé: and

Sophie Valcke 23/11/y 17:24

Supprimé: that

Sophie Valcke 23/11/y 17:24

Supprimé: might

Sophie Valcke 23/11/y 17:25

Supprimé: could

Sophie Valcke 23/11/y 17:25

Supprimé: of them