A DTC Proposal for the NGGPS Program Office
Global Model Test Bed

1 July 2016 – 30 June 2017

Introduction

The overarching objective for the National Oceanic and Atmospheric Administration (NOAA) Research to Operations (R2O) Initiative is the design, development, and implementation of the Next-Generation Global Prediction System (NGGPS), with a goal to provide enhanced forecasting capability from a few hours to 30 days. To solicit the participation of the science community, the NOAA R2O initiative includes the distribution of annual external announcement of opportunities (AO), as well as internal AOs targeted at NOAA and Navy laboratories.

To ensure these funded community efforts contribute directly to the development and implementation of NGGPS, an effective R2O process is needed. In particular, the research community needs an infrastructure that will enable it to efficiently and effectively use the NGGPS modeling system and/or its component models for research, which would, in turn, contribute to further improvement of NGGPS and its component models. To facilitate the R2O process for the continued development of NGGPS, the DTC proposed in FY15 a multi-year effort to engage the global modeling community by working with the developers to establish the infrastructure necessary to efficiently and effectively support distributed development and transition of the most promising innovations into operations. For year 1, the Global Model Test Bed (GMTB) focused its efforts on the physical parameterization aspect of the atmospheric model and sea ice modeling. For year 2, the DTC is proposing to continue its efforts in these areas. The proposed work is organized into the following five areas: 1) development and support of the Common Community Physics Package (CCPP) and Interoperable Physics Driver (IPD), 2) development and maintenance of physics testbed infrastructure, 3) testing and evaluation of atmospheric physical parameterizations, 4) community sea ice modeling, and 5) program management. All of these activities will be conducted through a collaboration with ESRL's Global Systems Division (GSD). This Statement of Work (SOW) represents only the NCAR portion of this effort.

Statement of Work

Development and support of the Common Community Physics Package (CCPP) and Interoperable Physics Driver (IPD)

Motivation

Parameterization of physical processes is an essential component of a numerical weather prediction (NWP) model. Ongoing development and improvement of physical parameterizations (PPs) is a critical part of continued global model improvement. Increased forecast skill can be expected as more physical processes are accounted for, and with the increased sophistication appropriate for the model resolution and application. Establishing an effective R2O process to ensure that physics development efforts – funded by NGGPS or by other sources – contribute directly to the continued development and improvement of NGGPS is a significant challenge. A modular physics suite accessible both in-line as part of a prediction model, and off-line for isolated testing, will enable physics innovation and contribution from the broader community. In response to this need, the DTC and NGGPS put forth the concept of a Common Community Physics Package (CCPP) coupled with an Interoperable Physics Driver (IPD) that would be setup and maintained by GMTB to facilitate efficient and effective collaborative development of next generation physics suites.

Project Description

During year 1, the GMTB made significant progress towards assembling the important components needed for establishing an IPD/CCPP package that will meet NGGPS's needs. The GMTB worked with EMC and the NUOPC Physics Interoperability group to assemble a document that lays out requirements for the IPD/CCPP software package and coding standards for the software. Work is underway to turn these requirements into a software design. For year 2, a software design for the IPD that includes the functionality necessary to meet the requirements developed in year 1 will be finalized in collaboration with our EMC partners. This design will be used to develop an incremental implementation plan for evolving the current IPD and GFS physics into an IPD/CCPP package that will minimize the impact of the IPD development on the on-going operational physics development. The design and implementation plan will also address any ways in which the CCPP will need to evolve to remain consistent with the IPD, and eventually meet the requirements put forth for the CCPP in year 1.

Establishing and maintaining an effective code management structure is critical for an effective R2O process. During year 1, the GMTB worked with EMC to develop an initial code management plan for the IPD/CCPP package. As the capabilities of the IPD/CCPP progress, the code management will need to evolve to stay current. To meet this need, the GMTB will refine and update the initial code management plan and define procedures to facilitate contributions to the IPD/CCPP from the broader community. This code management plan will include rules for updating the repository, and testing requirements for assuring updates to the repository do not have unintended consequences. Another key component of this plan will be establishing a Physics Review Committee that will be responsible for reviewing and approving changes to the physics package. Success of this process will depend on determining the appropriate membership of this committee and the protocols this group will follow.

A key component of any code management plan is the test suite necessary to insure the integrity of the evolving code base. Given the complexity of community contributions and operational requirements, rigorous software testing methodologies will be necessary to provide confidence in the software process. The GMTB will work with EMC to develop an initial test suite that will perform the necessary unit tests and regression testing for the combined IPD/CCPP package. As new physics parameterizations are approved and added to the CCPP, the test suite will be upgraded to include testing for the newly supported configurations.

Documentation is another key element of the infrastructure needed to facilitate distributed development. During year 1, the GMTB assembled most, if not all, of the pieces needed to document the initial IPD/CCPP capability, which is based on the operational GFS physics suite. As the IPD/CCPP package evolves, the documentation will also need to be updated to reflect the current capabilities. For year 2, the GMTB will refine the initial documentation by making changes to existing content to reflect corrections stemming from a scientific review by scheme authors and/or EMC physics experts. Documentation for year 1 focused on the individual physics schemes, whereas the interactions among the physical parameterizations within a physics suite are also important. Hence, during year 2, the GMTB will enhance the existing documentation by adding information on how the schemes depend on each other (i.e., what information is passed between the schemes in the GFS physics suite). A list of tunable and non-tunable parameters will also be added to the documentation for each scheme with physical descriptions (where possible), acceptable numerical ranges, and references to any sensitivity studies. As new physics suites are added to the IPD/CCPP package, it will be important to expand the documentation to include other suites. The GMTB will work with the developers of these suites to produce documentation following the template used to produce the documentation of the operational GFS physics suite.

To facilitate an efficient and effective R2O process, it will be important for physics developers to have a good understanding of the steps required to add a new physics package/suite to the IPD/CCPP. To meet this need, the GMTB will work towards providing overarching documentation that will provide information about how code for a physics parameterization/suite can interface with the IPD, the procedure for modifying the IPD to work with the new scheme and how to use the physics testbed described below to test the functionality of the new scheme within the IPD/CCPP framework. This documentation will also include information for developers who want to work directly with the CCPP code repository on the proper procedures for channeling development into the code repository. The overarching documentation will likely take the form of an online tutorial that can be used by developers on both "sides" of the driver (dynamic core and physics). This tutorial would reference the IPD and CCPP documentation as appropriate. In addition to this documentation, the GMTB will provide support to physics developers to ensure that all developers use the same reference code and understand the proper procedure for channeling their development into the CCPP.

Project Deliverables

- Software design that meets the requirements for the IPD/CCPP concept
- Implementation plan for IPD/CCPP capability
- Initial IPD/CCPP capability
- Code management plan and testing suite that are consistent with current IPD/CCPP capability
- Refined/enhanced documentation for GFS operational physics suite within CCPP

Development and maintenance of the physics testbed infrastructure

Motivation

To facilitate physical parameterization development, it is critical that the testing and evaluation performed by each developer be as uniform as possible. Infrastructure functions, including scripting, code management, diagnostic tools, offline component simulators such as a single-column model (SCM) and verification data sets, must be supported to provide a uniform "test harness" to enable in-depth investigation of various physical parameterizations. The test harness envisioned for NGGPS mimics the logical progression for testing newly developed parameterizations that typically takes place within the scientific community. Components are gradually added as one moves through the hierarchy until the full forecast model complexity is reached. It is designed to complement both the existing testing protocol at EMC and independent testing typically performed by parameterization developers. Figure 1 illustrates the hierarchical tiers of the test harness and represents how the GMTB envisions the division of effort (GMTB's likely role denoted by blue) and how the harness fits within EMC's existing testing framework.

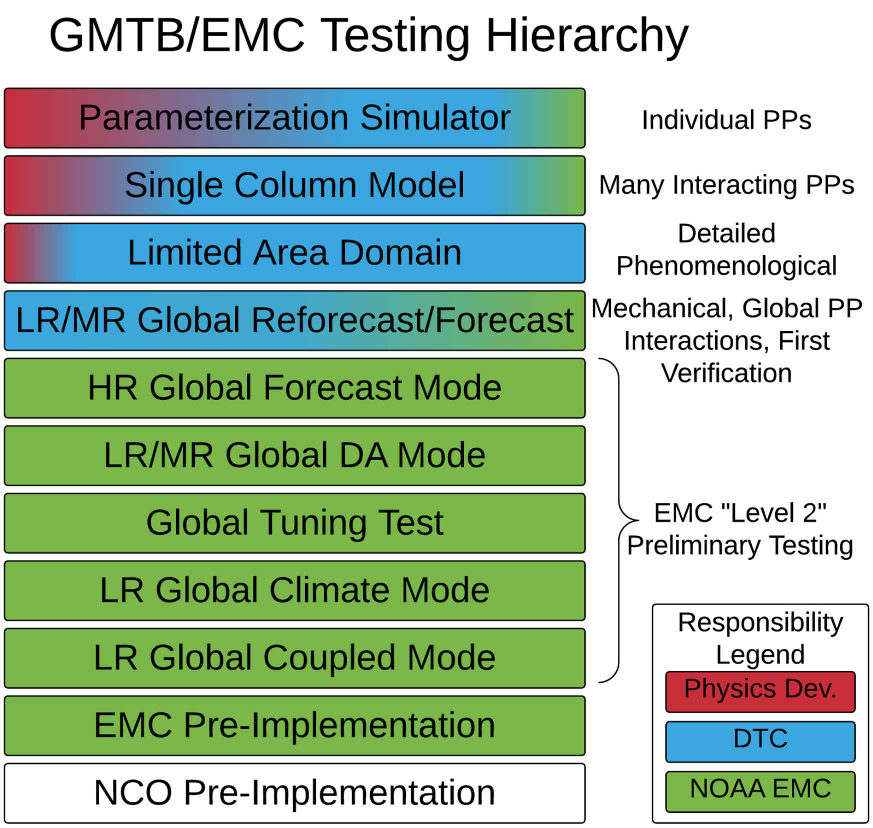

Figure 1. Diagram illustrating the testing hierarchy plan to support physics development for NGGPS. LR indicates low resolution, MR medium resolution, and HR high resolution. Color shading indicates where the different groups are anticipated to focus their efforts (red – physics developers, blue – GMTB task within the DTC, and green – EMC). PP stands for physics parameterization.

Project Description

In year 1, the GMTB established an initial physics testbed capability that includes two key capabilities within the logical progression for testing newly developed physical parameterizations: a Single Column Model (SCM) capability that builds on the IPD/CCPP package and a workflow for running the atmospheric component of the GFS without data assimilation and generating diagnostic information about the interactions of the physics packages and feedback on the large-scale flow (i.e., LR/MR Global

Reforecast/Forecast Mode). For year 2, the GMTB is proposing to enhance both aspects of this initial physics testbed capability, to begin to add a parameterization simulator capability, and to work with EMC to develop a plan for adding a Limited Area Domain or Cloud Resolving Model (CRM) capability in year 3. As the capabilities of this test harness progress and the number of developers who engage in this testing process grows, it will be important to make sure adequate computing resources are available for the GMTB and the physics developers to conduct extensive testing in order to assure the R2O process proceeds at the rate necessary to meet the NGGPS goals.

Single Column Model

The SCM capability established in year 1 provides the capability to run the operational GFS physics suite (the initial CCPP suite), using the current IPD, for idealized cases based on intensive observational periods. Continued development for year 2 will focus on remaining up to date with changes in the underlying IPD as functionality is added, and adding the capability to initialize and force the SCM with global model (e.g., GFS) output.

While running idealized cases following the GEWEX format is valuable, having the capability to run a SCM for any point on the global grid for any time period allows an investigator to better isolate and interrogate issues that appear during the full global testing. In addition, it allows for expansion of the atmospheric conditions available for focused testing (which are limited by the idealized cases available if effort is not made to create new idealized tests), and more concretely links testing within the hierarchies of the testbed.

As a longer-term project, and to provide an additional tool to be used with global physics tuning tests done at EMC, with a sufficient number of SCM test cases spanning a large variety of atmospheric conditions, the SCM may be used to explore a suite-wide parameter space. That is, if the tunable parameters for all schemes are externalized (and changeable via run-time scripts), the SCM can be run quickly for all test cases to rapidly test particular sets of tunable parameters. A similar approach was used by Golaz et al. (2007) to determine the best set of tunable parameters for the Cloud Layers Unified By Binormals (CLUBB) planetary boundary layer (PBL) and shallow convection scheme.

In the past, concerns have been raised about more-advanced parameterizations leading to degraded forecast performance as a result of physics suites being tuned to compensate for errors caused by other parameterizations, or even dynamic core numerical errors. In this framework, parameters may be set to values outside of their physically plausible ranges during the tuning process in the quest for the best overall model performance. The process described above using the SCM could potentially explore the tunable parameter space far more efficiently than using global runs (to help identify a suite-wide parameter set that works well with a more advanced physics scheme), and to shed light on why some tunable parameters provide better performance, with values set outside of their physically plausible range.

Workflow for Low/Medium Resolution Global Reforecast/Forecast Mode

To continue to facilitate the development and testing of operational and/or new physics parameterizations, the GMTB will expand the current inventory of diagnostic and verification tools to include additional capabilities. The GMTB will be responsive to user needs and be amenable to community input regarding additional tools to include in the testbed. Expanded capabilities will be tailored to specific research focus areas to help demonstrate the strengths and weaknesses of the physics parameterizations being tested. One such example of a key focus area for physics developers is cloud and radiative properties; to address this focus area, the GMTB will provide observation datasets as well as necessary scripts to diagnose and verify cloud and radiation fields (e.g., brightness temperatures and radiative fluxes). Example data sources include Surface Radiation Budget Network (SURFRAD) data and Geostationary Operational Environmental Satellite (GOES) data. As the need emerges, the GMTB may provide additional types of non-standard observation data sets via the testbed (e.g., data from field campaigns).

Another focus area will be to assist physics developers with testing physics parameterizations at higher resolutions. While the current capabilities are focused on medium-to-coarse resolution modeling, the GMTB will include tools to analyze, diagnose, and verify model output at higher resolutions. In year 2, the GMTB will be proactive forming collaborations with physics developers, in particular NGGPS Principle Investigators, in order to promote the use of the physics testbed. Having a facility to host a spectrum of evaluation tools will allow more effective and efficient physics development, community engagement, and provide an infrastructure that supports R2O and O2R.

The global workflows currently supported by the GMTB have the capacity to run a global model, post-processor, and a suite of diagnostics and verification. Due to the required architecture of select components (i.e., global model and post-processor), the workflow is only supported on Theia (NOAA research and development supercomputer). For some physics developers, compiling the necessary components and setting up scripts to run a global model (and necessary diagnostics and verification) can be time-consuming and resource intensive; it would be a significant benefit of the testbed to reach a larger audience by making the system more flexible and potentially portable.

EMC's Global Team has started to establish a global unified workflow system based on the NWP Information Technology Experiment (NITE) principles. The goal is to refactor the current set of scripts in order to have a customizable workflow that can support various systems for running experiments as well as easily repeating or replicating previous experiments. The GMTB will continue its dialogue and collaboration with EMC's Global Team and adopt new tools as the unified workflow system matures. It should be noted that while the pieces of the workflow regarding the model and post-processor are currently only supported on Theia, the diagnostics and verification components of the workflow were built with a broad user community in mind. This flexibility and portability allows for the diagnostics and verification scripts to be shared with a larger audience in a shortened time frame.

Parameterization Simulator

For year 2, it will be important to develop the capability to drive all parameterization types as a simulator. Much like a SCM can help to isolate physics suite errors from a fully interactive three-dimensional model, a parameterization simulator can help isolate errors at the individual parameterization level. In a parameterization simulator, developers working on each type of physical parameterization test their schemes by driving it with only the data that it needs to function. For example, a surface layer simulator would be driven by the physical surface properties and lower boundary layer values of wind, temperature, and moisture. The parameterization then calculates surface fluxes, given the input forcing, and the output can be compared with observations, or other schemes driven by the same data. For PBL schemes, the SCM can be considered a parameterization simulator if all of the other schemes are switched off and their contributions to the change in column properties are accounted for in the forcing. Some PBL-centric Global Energy and Water cycle Exchange (GEWEX) cases are formulated like this -- surface fluxes are prescribed, radiative heating rates are prescribed, and precipitation processes are switched off. While it is fairly clear how to do this for surface layer and PBL schemes, it is less clear how simulator capability would be accomplished for microphysics, radiation, and gravity-wave drag schemes. Input from community scientists will be valuable for adding this capability to the testbed, including organizing datasets to drive all simulator types.

Limited Area Domain

The limited area domain tier of the testing hierarchy is designed to test a physics suite as part of a three-dimensional model run with a regional domain. These tests afford the developer a relatively economical test of a physics parameterization or suite in a three-dimensional setting and are especially important for components of a physics suite that are affected by horizontal advection and vertical motion. To assure a seamless progression within the testing hierarchy, the choice of limited area domain capability should take into consideration the dynamic core selected for NGGPS. Given this decision is anticipated to be made by fall of 2016, year 2 of GMTB is an ideal time frame for GTMB to begin working with EMC to develop plans for the limited area domain tier of the testing hierarchy. The planning process will take into consideration the properties of the future global model system, as well as the potential to leverage existing capabilities.

Project Deliverables

- Enhanced SCM capabilities
- Initial parameterization simulator capability
- Enhanced workflow capability for low/medium resolution global reforecast/forecast mode

Testing and evaluation of atmospheric physical parameterizations

Motivation

To assure the test harness is used by the NGGPS-funded physics developers, as well as developers funded by other sponsors, it will be important for the GMTB to actively engage key developers in the application of the tools assembled by the GMTB, as well as the GMTB conducting its own T&E activities in consultation with EMC. Based on the DTC's experience with its Mesoscale Model Evaluation Testbed (MMET), the DTC is fully aware that simply building a testing infrastructure does not guarantee the community will make use of the tool.

Project Description

Following the first year's primary task of developing the tools and infrastructure for evaluating physical parameterizations, efforts for year 2 will focus on both continued development of the testbed tools and demonstration of the testbed's capability. The extent of the demonstration will largely depend on continued IPD development, and on which physical parameterizations are working seamlessly with the IPD. The initial IPD/CCPP development focused on assuring the NGGPS dynamic core candidates could interface with the current GFS operational physics suite. While the current CCPP includes code for schemes associated with past operational configurations and select advanced physics schemes that EMC is considering for future implementations, the communication between the IPD and these schemes needs to be checked to make sure everything is working properly.

Work is ongoing outside of the GMTB to develop members of an "advanced" GFS physics suite. GMTB will work closely with EMC and NGGPS Physics Team to assess which schemes are mature for testing in year 2. Possible candidates are, among others, the Simplified High-Order Closure (SHOC) PBL scheme, the Shikira-Sugyama cumulus parameterization, and the physics suite used in the High-Resolution Rapid Refresh (HRRR) model. What physics parameterizations are actually tested by the GMTB in year 2 will be determined by what schemes are scientifically mature, ready from a software perspective that includes seamless connection with the IPD, and the priorities of EMC. Multiple schemes may be considered for the lower tiers of the testing hierarchy, and based on the outcome of these tests a subset elevated to higher tiers of testing. The extent of GMTB's involvement will likely also depend on the interests of the scheme developer. The GMTB's role could range from assisting the developer with using of the testing harness to actually running the tests and consulting with the developer on the outcome of the tests. Regardless of the extent of the GMTB's role in the testing, all testing will be conducted in close coordination with the developers and/or appropriate Subject Matter Experts, as well as EMC.

As an example, testing of the SHOC parameterization could involve utilization of the SCM as a simulator, where only the PBL portion of the SHOC scheme interacts with the forcing. It can also be tested in a full SCM environment, where the entire GFS operational physics suite is active with SHOC replacing the PBL scheme and passing subgrid-scale cloudiness information to the radiation and microphysics schemes. For these lower tiers of testing, the modified physics would be compared to both the GFS operational "baseline" output and whatever observations and synthetic observations are available on a case-by-case basis. Analysis would likely focus on sub-grid fluxes of heat, moisture, and momentum, cloud profiles, precipitation rates, and time-integrated tendencies from all active physics schemes. For global testing, depending on available resources, cold-start reforecasts spanning a month or longer would be run using both the operational GFS physics suite and the SHOC-modified suite and evaluated using the verification tools available in the Physics Testbed. If possible, it would be useful to have runs of various horizontal and vertical resolutions to study resolution dependence of the results. Where possible, evaluation at each tier would attempt to include a measure of suite computational efficiency in order to develop a sense of price/performance ratio.

Now that the GMTB has established a basic test harness capability, it will be important to make the community aware this capability is available. Based on the DTC's experience with MMET, it will be important to work towards a publication in the literature that demonstrates how useful this framework can be to the physics development process. Hence, the GMTB will make it a priority in year 2 to prepare a manuscript based on one or more of the testing activities that demonstrates how the Physics Testbed can be a useful tool to physics developers.

Project Deliverables

- Report(s) on physics testing activities
- Manuscript on Physics Testbed submitted to appropriate journal

Establishing a community sea ice model framework

Motivation

To keep pace with the NGGPS development, NOAA needs to develop a strategy for sea-ice modeling. In year 1, the DTC hosted a workshop that brought together representatives from groups pursuing research related to sea ice forecasting with the primary goal of providing a recommendation to NOAA's NGGPS Program Office on the choice of a sea ice model for inclusion in NCEP's Unified Global Coupled System (UGCS). Given the use of a community-contributed and supported model in UGCS was raised as a priority for model selection, participants recommended the tentative adoption of the Los Alamos sea ice model CICE, pending follow-up testing and addressing concerns raised regarding model governance and differences in staggering between the grids used in the UGCS ocean models and CICE. The possibility of forming a consortium for CICE governance and support, including DOE, Navy, NOAA, NSF and primary contributors to CICE, was put forth as a possible framework for ensuring that the mission needs of operational prediction agencies using CICE can be incorporated into development priorities. Discussions directed at forming such a consortium are currently underway, but many details will need to be nailed down over the coming months before this concept can become a reality.

Project Description

As NOAA and the NGGPS Program Office navigate the pathway to establishing a consortium for CICE governance and support, the GMTB will continue to provide support to the NGGPS Program Office. Understanding some of the political sensitivities between groups involved in these discussions will be important for NOAA to successfully navigate these negotiations. The GMTB will help with this aspect of negotiations by continuing to enlist the assistance of sea ice experts in NCAR's Climate and Global Dynamics (CGD) Laboratory. At the same time, the GMTB will also provide advice on establishing a sea ice model code management plan that will meet the future needs of NGGPS.

Project Deliverables

- Support to NGGPS Program Office for engaging the sea ice modeling community in the process of setting up a community sea ice modeling framework.

GMTB Project Management

Motivation

The management of the GMTB will be supported by the DTC Director's Office. The DTC Director is responsible for the overall coordination of GMTB activities and maintaining strong ties with the community. The DTC Assistant Director helps the DTC Director with this overall coordination. Due to the distributed nature of the GMTB, the Director must rely on staff at the respective institutions to oversee the staffing, budgets and reporting to assure accountability. The DTC Director and Assistant Director will maintain regular communication with the NGGPS Program Office, to ensure NGGPS Program Office guidance is received on the execution of GMTB tasks in a timely manner.

Project Description

The DTC Director's Office will be responsible for the planning, coordination, management, and reporting of GMTB activities funded by the NGGPS Program Office. This work includes the development of a detailed work plan and budgets, and communication and coordination with EMC and other community partners to ensure smooth execution of GMTB activities. Workshops and community meetings are also central to NGGPS fulfilling its vision for involving the broad community in establishing and continuing to advance a state of the art global prediction system. The DTC will support these community workshops and meetings at the request of the NGGPS Program Office. For year 2, a moderate amount of program reserve will be held at the DTC Director's Office to support the requests of the NGGPS Program Office to support community meetings as needs arise.

Project Deliverables

- Quarterly reports on overall progress of NGGPS-funded activities
- Meeting support, as needed

References

Golaz, Jean-Christophe, Vincent E. Larson, James A. Hansen, David P. Schanen, and Brian M. Griffin, 2007: Elucidating model inadequacies in a cloud parameterization by use of an ensemble-based calibration framework. *Mon. Wea. Rev.*, **135**, 4077-4096.