MODELING WORKFLOWS

A BDTF white paper with accompanying findings and recommendation

Ad Hoc Task Force on Big Data

November 3, 2017

A BDTF White Paper

Modeling Workflows

I. Executive Summary

Modeling is a critically important aspect of NASA research that enables better understanding of the Earth and space observations that NASA collects and also serves as a conduit to the creation of products that can be used for further analysis and synthesis. Greater progress in modeling Earth, planetary, solar and inter-stellar systems by NASA scientists and grantees, as well as researchers at other Federal agencies, can be achieved if the workflows that are deployed to develop, assess and experiment with the models are substantially improved – more efficient and easier to use – and if steps are taken to more tightly connect the domain scientists with the information technology experts who can help modernize workflows.

Despite its importance, the manner in which NASA and other federal agencies model Earth, Earth-Sun, planetary and interstellar systems, generally described herein as “modeling workflow”, is not keeping pace with the rapid advancement of information technology and high-performance computing, methodological innovations, or the rapid growth of data volumes and complexity. Modeling workflows are largely ad hoc and little changed from when they were first conceived decades ago. The modeling workflows in place today have reached the point where time between generations of Earth and space system models is measured in years. Scientists are encountering: increasing difficulty confronting models with new data sources; insufficient scalability to petascale computing and data volumes; lack of reproducibility; and performance limitations due to lack of data compression or brittleness of aging data formats. For example, models used in the Earth and space sciences are unable to achieve even as much as 5% of peak capability on current high-performance computing platforms, and some codes access less than 1% of peak. Furthermore, there is a mismatch between the preparation of scientists involved in modeling Earth and space systems and the requirement for information technology mastery in a fast-paced open-source software development environment.

Substantial, possibly enormous, improvements in modeling workflows can be achieved through prioritized investments in computing hardware, workflow design and software, and education and training. Workflows for modeling Earth and space systems must be re-designed using a systems approach, including attention to modularization, parallelization, and automation. More accessible and affordable computing and data storage capacities, a more modern algorithm and code base, and, eventually, a cultural acceptance of the importance of workflow development and management, are all critically needed. NASA should explore new methodologies that can more effectively isolate signal from noise and advise computational experiment design. Cultural change should be effected by organizing education and training workshops, scientific conferences and journal special collections that bring together domain scientists and information technology experts. Further study is needed to determine the efficacy and cost-effectiveness for NASA to deploy data-centric gateways, virtualized environments, and tools for sharing and automating workflows.

II. Introduction

Modeling is a major aspect of the Earth and space science research conducted by NASA and other Federal agencies. The development of numerical models of the Earth system, planetary systems or astrophysical systems is essential to the linkage of theory with observations. Furthermore, the optimal use of observations that are often quite expensive to obtain and maintain typically requires assimilation or other forms of objective analysis that involves numerical models.

In the Earth sciences, there are three types of modeling. First, there are models of the physical climate system that encode the behavior of the atmosphere, the oceans, the land surface, and the cryosphere as well their interactions with portions of the biosphere, notably the vegetation on the land surface. These models are typically used in Earth science research for data assimilation. Data assimilation is a critically important function in the design of space-based observing systems which employ observing system simulation experiments (OSSE) and in the objective analysis of multi-sensor data sets to formulate a four-dimensional estimate of the state of the Earth system. The assimilation modelers consume and produce *large volumes of data*, requiring considerable high-performance computing and knowledge about diverse observing systems, including their sampling and error characteristics.

Second, climate models are used for climate projection, which is done either in prediction mode solving the equations governing the physical evolution of the Earth system numerically as an initial value problem or in scenario mode solving the same equations as a boundary value problem with specified future concentrations of radiatively active gases and aerosols. Climate predictions are useful for decision making in climate-sensitive sectors such as agriculture, water and energy resources. Climate change scenarios are useful for policy guidance, both nationally and internationally. The output of climate models used in these ways are typically large volumes, but not excessively so. It is the *variety of the data* that is most challenging in this area. For example, the Earth System Grid Federation (ESGF[1]) handles a huge number of different models, different components, different quantities and different sampling and averaging intervals, all contributed by an international community of modelers exercising their models with the same common experimental protocol. The experience in the ESGF is that it takes considerable processing to get the data into a uniform form that can be used with commonly available tools and models. It is necessary to refactor all aspects, including data documentation, to make the process of generating and analyzing this type of data efficient.

Third, more complex Earth system models are being developed by inter-disciplinary communities with applications to ecosystems, hazards, human and animal population/migration, and various biogeochemical processes. The interdisciplinary modelers, who often work in small teams, use *large data volumes*, particularly when modeling problems require high spatial or temporal resolution. However, the data sets *can also be high variety*. For example, the inter-disciplinary modelers tend to work with Level 2 retrievals, which typically have higher spatial resolution than Level 3 data on regular grids. This introduces complications such as the variety of spatial footprints for different sensors. Also, because the work is by definition multi-disciplinary, multiple dissimilar datasets are often brought to bear in an Earth system model.

---

[1] https://esgf.llnl.gov/

In space physics, the situation is different. Rather than large community modeling efforts, there are many models but none are community models. Most require considerable expertise to run. The models are frequently modified significantly for each case. The simulations in space physics are usually first principle theoretical models. The data in space physics frequently are sparse so there is little data assimilation. (Note: the Community Coordinated Modeling Center provides a service to tailor highly simplified models for scientists but they do all the work for the user.)

Modeling pipelines in solar physics share similarities with the first type of Earth Science modeling, but within a physical system dominated by magnetohydrodynamics, radiative transfer and material properties of highly ionized plasmas. There is a strong emphasis on the last step in the workflow where the models with data are compared. Since there is no "ground truth" data to validate either, solar physics relies heavily on transforming model outputs using raytracing and radiative transfer codes from three dimensional physical properties into synthetic data (spectra, filtergrams, etc.). Similarly, sophisticated inversion techniques are applied to the data to derive physical properties. Both of these activities are both data-intensive and computationally expensive and come at the tail end of the modeling pipeline.

The advance of information systems is contributing to the diminishing efficacy of modeling workflows in several ways. In addition to sheer data volume, data variety is being driven both by scientific development and sensor technology. Sensors have become exponentially more pervasive – the estimate is that 50 billion devices will be on the internet (Internet of Things, or IOT) by 2020, and many of these will be, or will at least be capable of, providing environmental data.

On the other hand, computing and networking hardware are not keeping pace. The underlying latency, bandwidth and capacity characteristics of various types of memory and storage devices have not kept up with the continuing relentless growth of these data volumes. The growth of data volumes is also outstripping improvements in data infrastructure price performance, putting budgets for other activities at risk.

Policy is also playing a role. Well intended governmental mandates to retain and make available to the public these data torrents further constrain budgets and also limits the options of data managers trying to cope with the previous technically-driven points.

These issues of data volume and variety have reached a point where human interaction with data is prohibitively inefficient and automated processes and tools are required. However, a major barrier to progress is that modeling workflows aren't deemed by practitioners to be a design problem. Traditionally, workflows have been created by a slow accretion of software, typically based on scripts haphazardly modified by a succession of scientists not typically trained in modern software engineering methods. These scripts are typically undocumented, difficult to understand and adapt, and able to do only one thing.

In many ways, this reflects the situation that modeling faced before the scientific aspirations drove application complexity beyond a poorly defined but easily recognized limit at which further progress simply couldn't be made without an overall design. This juncture coincided with a more or less independent increase in the complexity of computer architectures e.g. parallelism, cache hierarchies, etc. The two coincident developments forced a change towards community models, better and more modular software practices, and modeling frameworks for coupling, parallel input/output (I/O), etc. This transformation in the practice of Earth and space science modeling is still underway.

Modeling workflows are just beginning to reach the same tipping point, being driven by a confluence of a variety of scientific and architectural concerns related to the increasing complexity and sophistication of the questions being asked of the data and the underlying capabilities of the available I/O subsystems and storage technologies.

III. Statement of the Problem

The workflows employed in modeling the Earth system, the Earth-Sun system or planetary and astrophysical systems in the 2010s differ very little in general from those used in the 1970s. Once a code is “frozen” – a process that often involves ad hoc testing and arbitrary criteria – the code is applied to either simulation or prediction research problems, often exercised against a pre-defined protocol determined by the research community. The computational campaign is conducted in batch mode on the largest available supercomputers, and a set of output variables are saved on long-term storage for future analysis. The selection of variables, resolution and frequency of saves for the output data sets is often an ad hoc process as well. The data accumulated in this way, are subsetted and copied to volatile media for analysis using a wide range of techniques including statistical and dynamical data models. While a number of tools have been created over the past several decades to aid one or another of the steps in this workflow, the overall workflow has received relatively little attention and is now, with the volume of output reaching 10s to 100s of petabytes, far from optimal.

The workflows employed to process these large volume and high variety data sets are typically ad hoc collection of scripts that were not designed for the purpose but accumulated over a period of time to solve sub-problems. As a result, there are several problems that cannot be properly handled by existing modeling workflows.

The ability to "onboard" new datasets into models. This involves (a) understanding the dataset content, sampling, averaging and quality, and (b) suitably preparing the data for use without losing or significantly altering the information content.

For some workflows, the massive data volume is in itself a problem. Data production rates are being driven by advances in underlying information system technologies, which continue to produce exponentially faster computations. In general, modeling data production rates tend to track the improvements in sustained computing capacity, which continues to double every couple of years. Likewise, improvements in the capabilities of observing systems, especially on space-based platforms, are being driven by the underlying digital technology as well, producing a similar exponential curve in their data generation.

Reproducibility. Climate and weather is sensitively dependent on initial conditions. This has led to a belief that computational experiments and the data they produce should be absolutely reproducible at a “bit-for-bit” level – at least for model development purposes – so workflows are created with the intention of maintaining reproducibility. Unfortunately, reproducibility is at best an ephemeral property, and may not exist in a practical sense at all. It is ephemeral, because compiler and system software changes alter the arithmetic and the order of the arithmetic being performed, and machine architectures executing the arithmetic may become obsolete. For example, global reductions (as found in iterative solvers) that do not preserve the order of addition for performance reasons, will get different answers from run to run. Reproducibility may not be possible for the more fundamental reason that transistors on the scale of 10 nm can be affected by the environment. An experiment at Oak Ridge on the *Jaguar* system a few years ago showed that fast neutrons, caused by cosmic rays impacting the Earth's atmosphere, were causing bit flips in *Jaguar's* memory at a rate of several hundred per second[2]. Error correction code in the system was catching these, which is why the system ran at all. But the probability of a double bit flip that cannot be corrected and doesn't cause a program to stop is non-zero. Since exactly the same simulation is rarely rerun, no one really knows how often this happens in practice.

It is quite possible that reproducibility at the "bit-for-bit" level is a constraint that is no longer affordable, and also one that will be increasingly hard to retain in future systems. An example from the computational realm: the use of FMA (Fused Mult-Add) in computations can introduce an ambiguity in the order of addition and multiplication that affects reproducibility. However, not using FMA in order to satisfy the reproducibility constraint means that performance takes a hit, energy consumed per simulation increases, and scientific progress is limited.

Data compression. It has been demonstrated that compression algorithms exist that can save as much as 80% of the data storage requirements of model output, termed "lossy" compression, without altering the statistical properties of the overall data set. Nevertheless, current workflows enforce loss-less compression at the expense of memory and network bandwidth consumption and unnecessary data storage.

Inherent performance limitations of standard data storage formats. Self-describing file formats became standard within the past two decades. However, many model "history" files, for example those employing unstructured grids, have a large fraction of their storage footprint taken up by the metadata descriptors needed to document the file contents. Nonsensically, these large headers are repeated over and over in each file, creating significant storage overhead. Workflows that employ a more linked object approach would reduce these types of overheads.

These factors, taken collectively, are creating an untenable situation for those performing and supporting Earth system and space science.

In stellar astrophysics, the main impediment is the availability of resources to hire the types of professionals needed to achieve continuing development in a fast-moving, open source environment. Ensuring the education and engagement of the next generation of developers is critically important, but it is challenging for them to put in substantial time for development given the funding climate and the still unfortunate lack of recognition in the science community of the value of building such a remarkably impactful instrument that many get to use.

IV. Examples of Approaches

---

[2] V. Sridharan and D. Liberty. A study of DRAM failures in the field. In *Int. Conf. on High Performance Computing, Networking, Storage and Analysis (SC)*, 76:1–76:11, 2012.

1. *Magnetospheric models*. There are large-scale system models that mostly are magnetohydrodynamic (MHD) simulations. In these models, the solar wind provides the input to the simulation and the output data are fluid and electric and magnetic field parameters in the magnetosphere. They use either ideal solar wind input or input derived from solar wind observations. They are fully self-consistent. For simulations driven by solar wind data comparison with observations is possible. There is very little data assimilation because the data are too sparse. There also are purely theoretical models that are used to study local or quasi-local phenomena. These are either hybrid (fluid electrons and particle ions in addition to the fields) or kinetic (particle ions and electrons plus the fields). They require vast resources in terms of computer time and produce "large data". The particles and fields in these models also are self-consistent. The initial and boundary conditions are ideal. Very recently, coupled MHD and kinetic (particle-in-cell or PIC) simulations have been developed. Here the initial and boundary conditions are set by the solar wind for the MHD calculation, while the MHD calculation sets the initial and boundary conditions for the PIC code.

Even though the models seem very different, the workflow is not that different especially after the first step. The researcher modifies the code, sets the boundary conditions and runs the simulation. Then the output or a subset is brought back to the researcher's home institution and analyzed. The output is a time series in 3D space. Because there is so much data (one example combined MHD and PIC code produces about 0.5 TB per time step and runs ~10^5 time steps), the return of even a subset takes days or more.

There are ways to make this faster by exploiting the latest developments in networking. The Pacific Research Platform (Pacific Wave[3], CalREN HPR[4], CENIC[5]) provides high speed (>100 Gb/s) communication links between research labs and universities along the pacific rim. There are vast underutilized resources available. Tests of the technologies currently underway between UCLA and Ames have found that a limiting factor is the rate at which the user's system can write to local storage. A project at UCSD (Smarr, presentation to the BDTF) is developing technologies to increase the writing speed to local systems.

2. *Earth system models*. It has long been recognized that Earth system modeling involves an elaborate process of revising models, running them against standard protocols, and checking to see if the revision has the desired effect and/or undesirable side effects. Because Earth system models are multi-component (atmosphere, ocean, land surface, and cryosphere) and multi-physics (dynamical governing equations, representation of radiation and other forms of energy transfer, and bio-chemical transformations), the improvement phase is typically applied separately for different components or different physical processes, and the running/checking phases are sub-cycled for the separate components/physics. The running/checking loop is then applied to the full Earth system model. The standard run protocol typically involved many hundreds to thousands of years of simulation, either sequentially to reach equilibrium after a given revision is applied or in parallel to an ensemble of simulations that are equi-probable by construction in order to evaluate the robustness of results. The checking phase may require the interrogation of very large data sets,

---

[3] http://pacificwave.net/

[4] calREN _- California Research and Education Network – CENIC's network: https://cenic.org/network/network-maps

[5] CENIC - Corporation for Education Network Initiatives in California: https://cenic.org/

e.g., current simulations amount to 100s of terabytes or petabytes, so that analysis operations must be able to achieve million-fold data reduction.

This model development cycle is unacceptably long – at present, large Earth system models are updated to new full versions about every 5-7 years, and even then, because of the high level of complexity of the code, many branches of the data flow are untested. Several groups are working on ways to streamline this cycle by applying parallelism wherever possible. The model codes are being refactored for accelerator-based systems and, in the process, the algorithms are being evaluated for potential concurrency of operations. Similarly, the data interrogation phase is being rewritten to exploit parallel analytic methods and automated data reduction techniques. In some cases, notably at NCAR, systems with high-speed disk and ample computational capacity are being deployed solely for the data analytics step in order to provide a data-centric science gateway. Parallel visualization and feature recognition tools are being deployed to speed up the checking phase.

Another trend in Earth system modeling is the development of community models. The two principal such examples are the Community Earth System Model[6], hosted by NCAR and used primarily by scientists in the U.S., and EC-Earth[7], built around the core atmospheric model from the European Centre for Medium-range Weather Forecasts, which is used by 10 countries in Europe. In both cases, there is a support system for users, considerable documentation of the code, configurations, data sets and ancillary information, and a very broad community of users who conduct experiments with the models and analyze the output of standard-protocol simulations.

3. *Stellar astrophysics*. In stellar astrophysics, an open source code called MESA (Modules for Experiments in Stellar Astrophysics) is used for all of stellar structure and evolution. The goal of the MESA development team has been to create and continuously develop a tool that allows for scientists from all over the world (now about 1000) to use the tool themselves to explore the properties and behaviors of stars. It's been used to model everything from hot Jupiters around nearby stars to the collapse of massive stars to form black holes. It is estimated that100's to 1000's of MESA runs with ~12 threads/run are needed (a big data problem in the higher volume and more variety mode).

V. Recommended Approach(es)

Improving modeling workflows to address the problems described above requires a systems approach with balanced investments in hardware, software and human resources. It is hard to estimate where the emphasis should go because the system is coupled: for example, investments in new architectures and software may change workflows in ways that will in turn affect data retention rates, which will in turn affect the relative investments required between storage and analysis systems. The approach includes these steps:

1. *Hardware investments*. Initially, there are infrastructural design issues associated with emerging storage technologies that will require mid-range investments in prototype systems and system software for testing and experimentation. HPC and data storage investments may require sizeable increments in order to keep up with data volumes.

---

[6] http://www.cesm.ucar.edu/

[7] http://journals.ametsoc.org/doi/pdf/10.1175/2010BAMS2877.1

2. *Software investments*. Substantial investments in workflow software modernization efforts are needed, which should be coupled with hardware investments. Some examples are given below.
3. *Accessible and affordable compute and storage resources*. Some workflows have educational value to students in the classroom, so easy availability of compute resources on demand for web-enabled computational exploration is a growing need. Some current options, such as supercomputer hosting and commercial cloud computing services may not easily suit this need in the long term.
4. *Experimental design methodologies*. Many modeling protocols have the general nature of explorations of a parameter space; however, insufficient attention is typically paid to the design of such exploration, which can be computationally very expensive, and may involve the interpolation of calculated results in sparsely sampled parameter space, as well as the statistical fitting of the computational results to observational data. One example of improved sampling of parameter space is the use of irregular grids called 'Latin squares' that can achieve higher accuracy with fewer runs. Considerable recent mathematical effort has been made specifically to quantitatively optimize the design of computer experiments[8], and improved designs can be obtained quite simply using public codes[9] - see Methodology white paper.
5. Scientific investments should focus on developing new data analysis paradigms using the latest ideas in spatial statistics and machine learning. This will require an interdisciplinary approach.
6. Some investment in institutional culture change, consensus building and standards development may be required. A suite of activities, including training workshops, scientific conferences and journal special collections are needed to ease the scientific communities into more modern practices.

In order to prioritize investments, it is necessary to target critical path issues to accelerate modeling workflows. Having a specific project objective and timeframe is helpful to set requirements and drive development of workflow improvements. The main critical path issue is time to solution. That is, the modeling workflows in place today have reached the point where time between generations of Earth and space system models is measured in years. There are several bottlenecks that are slowing progress, not least being the efficiency of codes on current high-performance computing platforms. Many Earth and space system models are running at less than 5% of peak performance as measured against optimized linear algebra benchmarks. Some codes are achieving less than 1% of peak. Nevertheless, it is difficult to identify the single most important factor contributing to the downward trend, so a systems approach is needed that examines the full cycle of model development, confrontation of models with observations, and analysis. Some examples of work that could be undertaken to accelerate workflows include:

- Modularization of modeling workflows, through the identification and creation of reusable workflow primitives
- Parallelization and optimization of these workflows and primitives to achieve greater unit efficiency
- Automation of multistep workflow execution, e.g., employing CYLC, a tool used at a growing number of modeling centers

---

[8] e.g., arxiv.org:1601.05887

[9] e.g., R/CRAN: https://www.r-project.org/ and https://cran.r-project.org/
Lhs (Latin Hypercube Sampling):
http://finzi.psych.upenn.edu/R/library/DoE.wrapper/html/lhs.design.html
http://finzi.psych.upenn.edu/R/library/MaxPro/html/MaxPro-package.html

Longer range investments on more innovative capabilities should be considered, such as:

- Virtualized environments for data analysis
- Improved scalability, including memory scalability for data assimilation systems
- Investigating tolerance for lossy data compression to reduce data flow and modern signal processing methods for increasing signal detection and reconstruction (e.g. compressive sensing[10])
- Developing a new generation of data-centric systems that incorporate novel memory and storage technologies (e.g. nonvolatile and NAND-base), and storage hierarchy management tools
- Tools and community mechanisms for the vetting (peer reviewing), preserving and sharing of workflows
- Automation of workflow composition

Beyond these innovations, continued and enhanced investments should be made in the fundamental statistical and machine learning algorithms that will be needed to obtain new insights from the data. Parallelism and performance optimization should be designed into these new approaches from their inception through a strongly interdisciplinary focus spanning the computational science, computer science and spatial statistics communities.

Additionally, there is a desire to make the data more *accessible*, in all senses of the word. That means: (1) easier to understand for people outside the discipline in question; (2) easier to obtain, for exactly the space-time-feature space desired; and (3) easier to use and incorporate into models, once obtained.

Earth system model output data can be made easier to obtain by some investments in software and "the cloud", e.g., in some cases by reorganizing high-value datasets to make them more tractable, especially as a function of time.

Making Earth system model data easier to understand and easier to use are mostly problems that can be addressed with software and human resources. The human resources is largely a matter of getting the community, including data producers, archives, and end users to work in concert to improve the consistency and usability of data product design, organization, access methods and documentation.

VI. Conclusions

This white paper outlines several aspects of modeling workflows that are summarized in the following findings and recommendations.

---

[10] Candès, Emmanuel J.; Romberg, Justin K.; Tao, Terence (2006). Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, **59** (8): 1207–1223.

Finding: The BDTF finds that workflows, used by NASA as well as other federal agencies to model Earth, Earth-Sun, planetary and interstellar systems, are largely ad hoc and little changed from when they were first conceived decades ago. This gives rise to:

- increasing difficulty confronting models with new data sources
- insufficient scalability to petascale data volumes
- lack of reproducibility
- performance limitations

Finding: The BDTF finds that there is a mismatch between the preparation of scientists involved in modeling Earth and space systems and the requirement for information technology mastery in a fast-paced open-source software development environment.

Recommendation: The BDTF recommends that NASA should make the necessary investments in computing and analysis hardware, workflow software and education and training to substantially accelerate modeling workflows. NASA should take the lead to make substantial increases in:

- accessible and affordable computing and data storage capacities
- software modernization
- resources to develop new data analysis paradigms
- education and training workshops, scientific conferences and journal special collections to effect a culture change.

Immediate efforts that can contribute to accelerating modeling workflows include:

- modularization
- identifying and implementing concurrency in workflows and algorithms
- automation of repetitive steps in modeling workflows.

Longer-range investments whose potential NASA should investigate include:

- virtualized environments
- research on memory and processing scalability
- lossy data compression and more advanced methods for signal detection
- data-centric science gateways
- platforms for sharing workflows
- automation of the creation of workflows

Workflows designed and implemented several decades ago are no longer capable of keeping pace with the growing complexity of the model development cycle, the exponentially growing volume of input, output and validation data, and the rapidly advancing computing environment for both high-performance computing and large-memory data analysis. If the above recommendation is not adopted, the outdated modeling workflows employed by NASA and other Federal agencies will waste valuable high-performance computing, data storage and human resources and will increasingly hamper progress.

Great progress in modeling Earth, planetary, and inter-stellar systems by NASA scientists and grantees can be made if the workflows that are deployed to develop, assess and experiment with the models are substantially improved and made more efficient and easier to use, and if steps are taken to more tightly connect the domain scientists with the information technology experts who can help modernize workflows and accelerate progress.

VII. Acknowledgements

The authors appreciate the inputs made by:

- Jeanne Behnke (NASA GSFC)
- Lars Bildsten (KITP)
- Rich Loft (NCAR)
- Chris Lynnes (NASA GSFC)
- Ray Walker (UCLA)

All their very useful thoughts are gratefully acknowledged.

VIII. Definitions and Acronyms

Definitions

Data Assimilation: The process by which observations are merged with the state of a numerical model in such a way as to minimize an estimate of error.

Earth System Model: The governing equations, the algorithm(s) used to solve them, or the computer code that instantiates the algorithm(s), or all three, that describe the time evolution of the Earth system, considered as an interactive ensemble of components including the atmosphere, land surface, oceans, sea ice, ice sheets and glaciers, biota influencing weather and climate, and chemical constituents cycling through all of the above components.

Earth System Modeling Workflow: The programmed and repeatable pattern of computational tasks, enabled by the systematic organization of high-performance computational, disk and tape resources along with efficient software, that numerically integrate and Earth System Model. The computational tasks include (1) the acquisition and pre-processing of input data (e.g. assimilation of observations to produce initial conditions and boundary conditions); (2) execution of the dynamical solver; (3) storage of relevant output data in volatile, short-term and long-term media; (4) retrieval of subsets of the output data for post-processing and analysis; and (5) analysis and interpretation of subsetted output data, including comparison with analogous model outputs and observations. The storage and retrieval of output data (tasks 3 and 4), can be consolidated within task 2 to generate concise post-processed output as the solution evolves. Often multiple generations of models must be rigorously and quantitatively compared, necessitating the long-term storage of output data, together with the comparable observational data, in compatible formats.

Acronyms

BDTF	Big Data Task Force
ECMWF	European Centre for Medium-range Weather Forecasts
GSFC	Goddard Space Flight Center
I/O	Input/Output
KITP	Kavli Institute for Theoretical Physics
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
UCLA	University of California at Los Angeles

BDTF Findings & Recommendation for Model Workflows

Background: See BDTF white paper Model Workflows.

Finding: The BDTF finds that workflows used by NASA as well as other federal agencies to model Earth, Earth-Sun, planetary and interstellar systems are largely ad hoc and little changed from when they were first conceived decades ago. Scientists are encountering:

- increasing difficulty confronting models with new data sources
- insufficient scalability to petascale data volumes
- lack of reproducibility
- performance limitations

Finding: The BDTF finds that there is a mismatch between the preparation of scientists involved in modeling Earth and space systems and the requirement for information technology mastery in a fast-paced open-source software development environment.

Recommendation: The BDTF recommends that NASA should make prioritized investments in computing and analysis hardware, workflow software and education and training to substantially accelerate modeling workflows. NASA should take the lead to make substantial increases in:

- accessible and affordable computing and data storage capacities
- software modernization
- resources to develop new data analysis paradigms
- education and training workshops, scientific conferences and journal special collections to effect a culture acceptance of the importance of workflow development and management.

Immediate efforts that can contribute to accelerating modeling workflows include:

- adopting a systems approach to designing workflows
- modularization
- identifying and implementing concurrency in workflows and algorithms
- automation of repetitive steps in modeling workflows.

Longer-range investments whose potential NASA should investigate include:

- virtualized environments
- research on memory and processing scalability
- lossy data compression and more advanced methods for signal detection
- data-centric science gateways
- platforms for sharing workflows
- automation of the creation of workflows.

Rationale: Workflows designed and implemented several decades ago are no longer capable of keeping pace with the growing complexity of the model development cycle, the exponentially growing volume of input, output and validation data, and the rapidly advancing computing environment for both high-performance computing and large-memory data analysis.

Consequences of not adopting the recommendation: The outdated modeling workflows employed by NASA will waste valuable high-performance computing, data storage and human resources and will increasingly hamper progress.