



parVis



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Update on progress (since Sept 30, 2011)

ParVis team

4th ParVis All-Hands meeting

March 22-23, 2012

Argonne National Laboratory

Progress made in all ParVis areas:

- Developing new Parallel Climate Analysis Library (**ParCAL**) and developing a parallel version of NCL (**ParNCL**) with it for DAV of ultra-large climate data sets.
 - Using DOE ASCR software such as MOAB (for grids) and Intrepid (for gradient operations) and PNetCDF (for parallel I/O).
- Speeding up current diagnostic workflows with Swift-based task parallelism.
- Taking advantage of hardware opportunities (GPU's, clouds) and avoiding limitations (disk space).

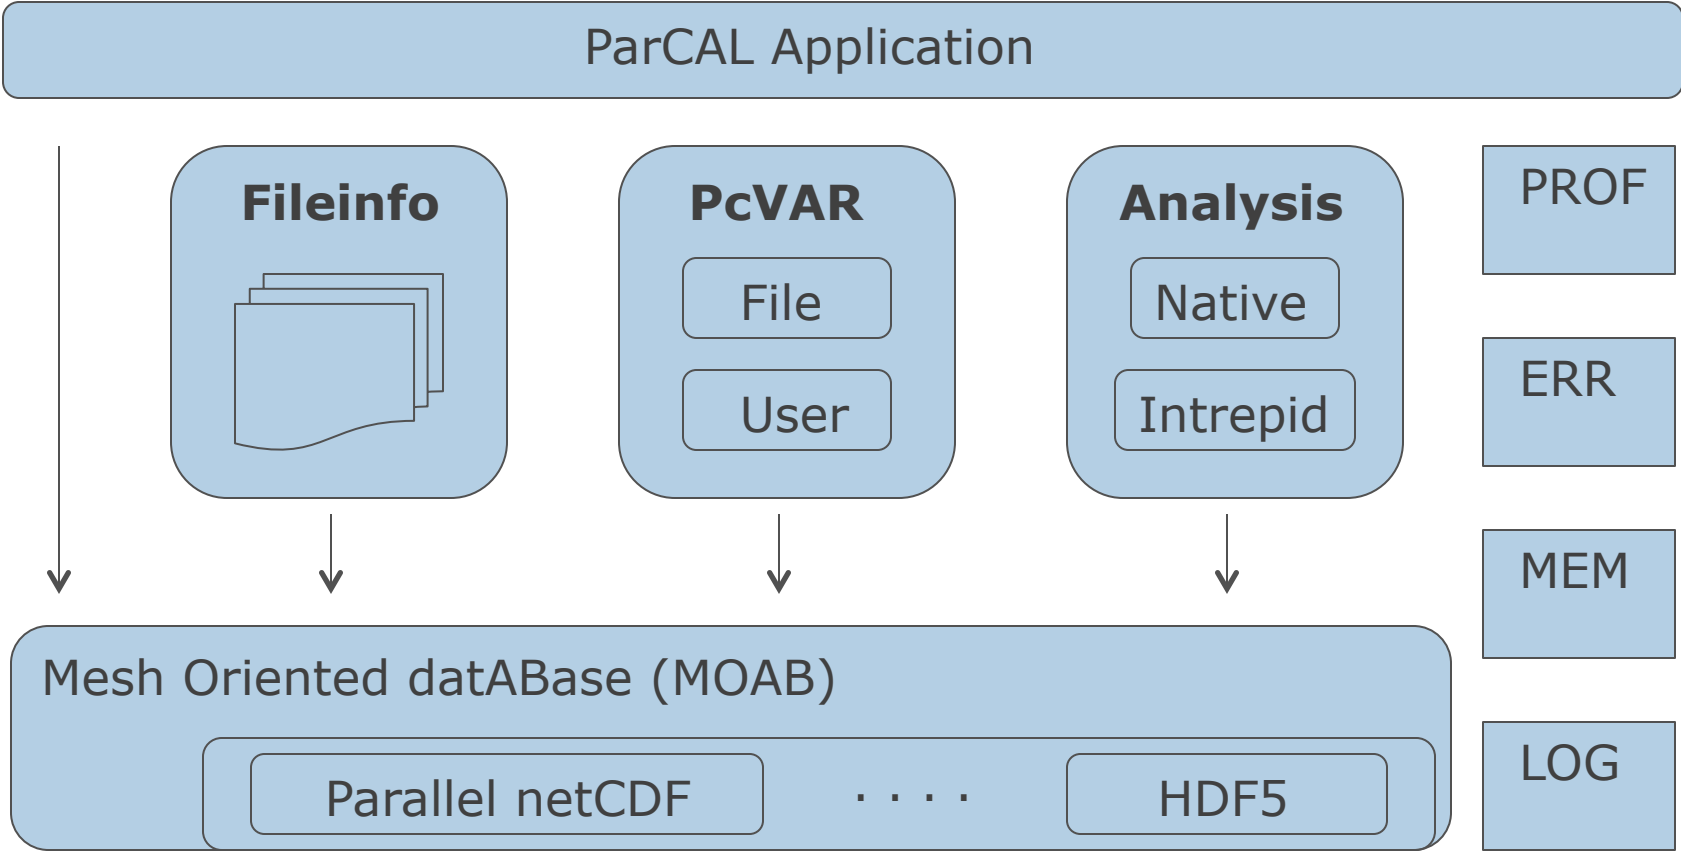


ParCAL: Parallel Climate Analysis Library

- Being built using PnetCDF, MOAB and Intrepid
 - MOAB and Intrepid have already solved the hard problem of how to represent and operate on structured and unstructured grids distributed over processors.
- Provides data-parallel core to perform typical climate post-processing currently done by either NCL or NCO.
- **Will be able to handle unstructured (CAM-SE (HOMME) and semi-structured grids in all operations (e.g. “zonal average”) by building ParCAL on MOAB and Intrepid. Will support parallel I/O by using PnetCDF.**



ParCAL Architecture



Development with Intrepid (a component of ParCAL)

- Divergence and vorticity
 - developed parallel versions using Epetra package from Trilinos
 - will update current ParCAL implementation with new tag_reduce functionality from MOAB
- Streamfunction and velocity potential
 - implemented with Intrepid for global velocity fields
 - investigating approach for limited domains
 - will incorporate into ParCAL
- Irrotational and non-divergent velocity components
 - implementation underway



Calculating Streamfunction and Velocity Potential with Intrepid

- The finite element method is used to solve the following weak equations for streamfunction and velocity potential using Intrepid

$$\int \nabla \psi \cdot \nabla \varphi \, d\Omega = \int \mathbf{v} \cdot (\mathbf{k} \times \nabla \varphi) \, d\Omega$$
$$\int \nabla \chi \cdot \nabla \varphi \, d\Omega = \int \mathbf{v} \cdot \nabla \varphi \, d\Omega$$

- Periodic boundary conditions along the latitudinal boundary and Neumann boundary conditions at the poles are used

$$\int_{\Gamma} \left(\frac{\partial \chi}{\partial n} - \mathbf{v} \cdot \mathbf{n} \right) d\Gamma = 0 \quad \int_{\Gamma} \left(\frac{\partial \psi}{\partial n} - \mathbf{v} \cdot \mathbf{t} \right) d\Gamma = 0$$

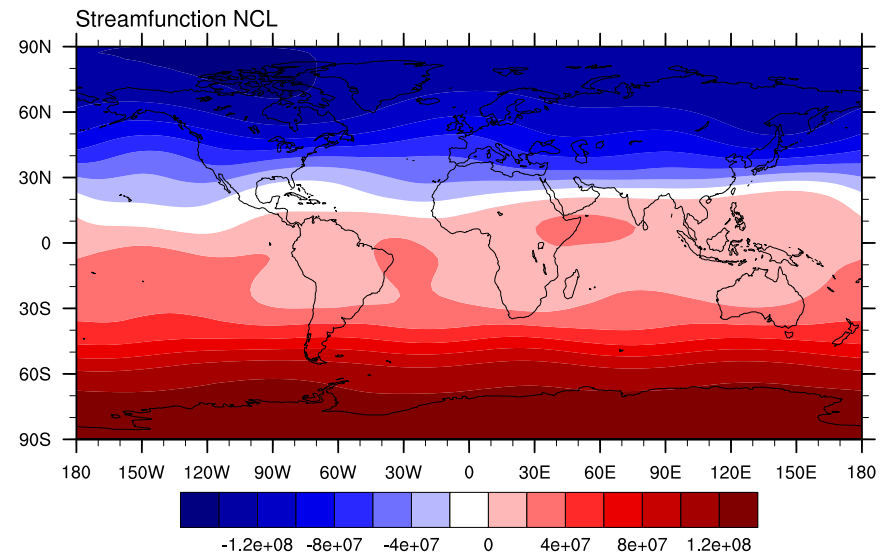
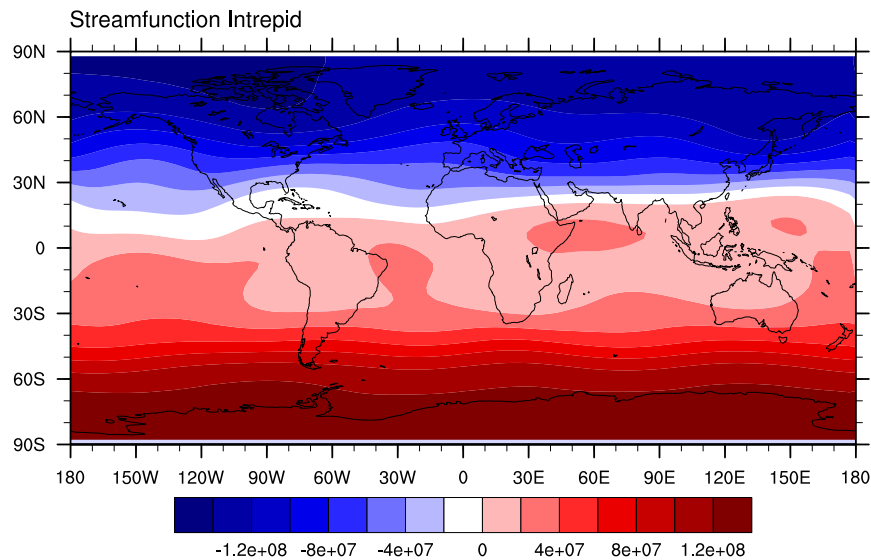
- The weak equations hold on arbitrary subdomains thereby enabling calculations from regional velocity data
- Intrepid can support solution of these equations on triangles and quads and eventually on polygons.



Calculating Streamfunction with Intrepid

Intrepid
finite element method

NCL (uv2sfvpG)
spherical harmonics



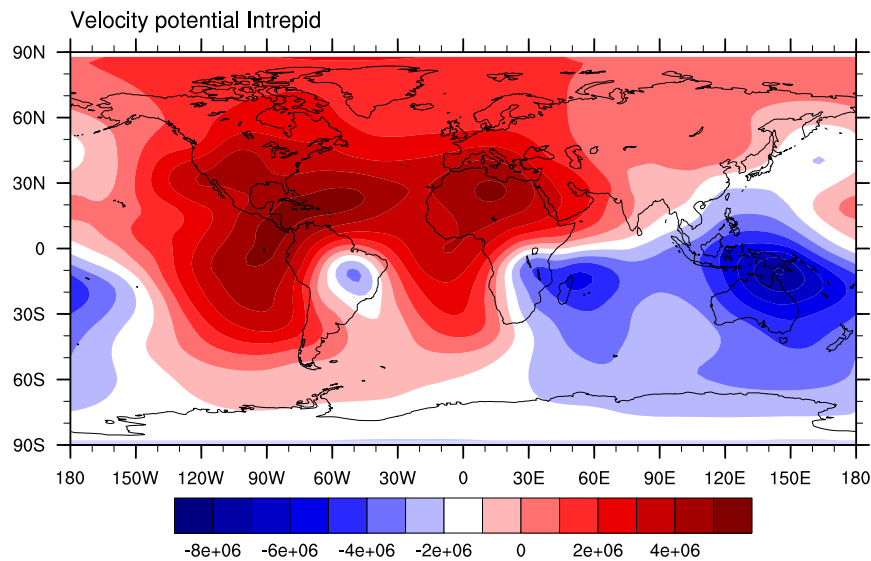
$$\nabla^2 \psi = \nabla \times \mathbf{v}$$



Calculating Velocity Potential with Intrepid


Intrepid

finite element method



NCL (uv2sfvpG)

spherical harmonics



$$\nabla^2 \chi = \nabla \cdot \mathbf{v}$$



ParCAL Development

- Integrated Intrepid-based algorithms into ParCAL
 - divergence, vorticity
- Updated algorithms to use MOAB partition method (SQIJ)
 - Gather (internal to ParCAL), NCL functions dim_avg_n, max, min
- Implemented new NCL algorithms
 - dim_max_n
 - dim_min_n
 - dim_median_n
- Miscellaneous items
 - Added tests of different dimensions (time, lev, lat, lon) for all *_n algorithms
 - Added more test configurations to nightly build/test system
 - Added support within MOAB and ParCAL to read global and variable attributes
 - Updated installation documentation and README
 - Removed all of the warnings generated by GNU g++ compiler

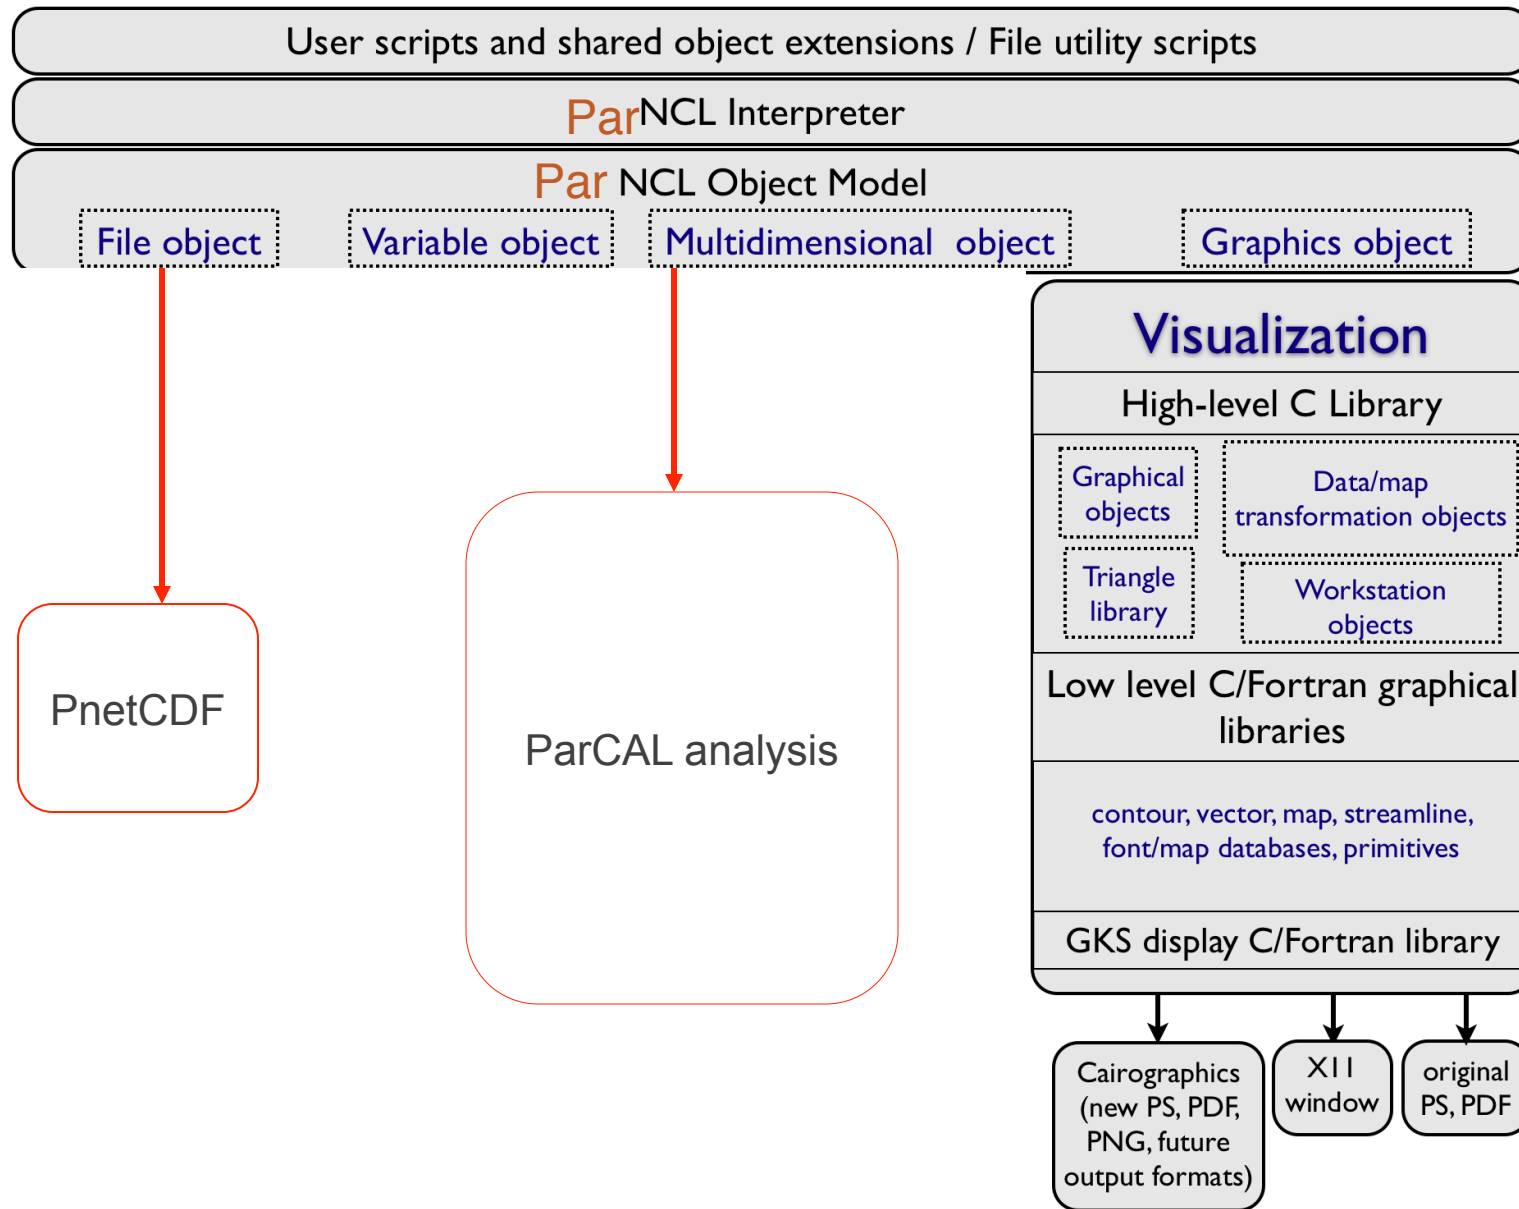


ParCAL Function Table

NCL function Group	NCL Functions	ParCAL Function
File IO	addfile, addfiles	fileinfo, pcvvar
Spherical Harmonic Routines	dv2uv* (4 funcs)	divergence
Meteorology	uv2dv_cfd	divergence
Spherical Harmonic Routines	uv2dv* (4 funcs)	divergence
Meteorology	uv2vr_cfd	vorticity
Spherical Harmonic Routines	uv2vr* (4 funcs)	vorticity
Spherical Harmonic Routines	uv2vrdrv* (4 funcs)	vorticity, divergence
General Applied Math	dim_avg, dim_avg_n	dim_avg_n
General Applied Math	dim_max, dim_max_n	dim_max_n
General Applied Math	dim_min, dim_min_n	dim_min_n
General Applied Math	dim_median, dim_median_n	dim_median_n
General Applied Math	max	max
General Applied Math	min	min
Variable Manipulators	delete	pcvvar
		gather



ParNCL architecture



ParNCL (Progress update)

- ParNCL supports `addfiles()`, NCL's multi-format file reader
 - Time slices of the variable read from file works
 - Need to test all corner cases
 - Only reads NetCDF for now (using PNetCDF).
- ParNCL supports calculating vorticity
 - Based on the current implementation in ParCAL
 - Specifically, duplicating function of `uv2vrG_Wrap()`
 - Adding support for other variations should be trivial from ParNCL perspective
- Simple math operations on distributed multidimensional variables work
 - `abs`, `cos`, `asin`, `atan*`, `cos`, `exp`, `fabs`, `floor`, `log`
 - Need to test more cases



ParNCL (Progress update)


- Addition, subtraction of distributed multidimensional data works
- Scaling a distributed multidimensional array by a scalar works
- Build changes
 - Can build serial and parallel versions separately
 - Can disable profiling layer at build time and runtime
 - ParNCL executable is now named “parncl”



NCL Version 6.1.0-beta update

Scheduled mid-to-late April 2012


- Main features:
 - Overhaul of display model to allow for 32-bit color (RGBA versus RGB), transparency, > 256 color maps
 - Significant improvements made to raster-drawing algorithm
 - Interface to Earth System Modeling Framework (ESMF) regridding capabilities will be included



parVis

New raster drawing features

- Rewrite of NCL raster plotting that promises a significant speed up (more than a factor of 40 in preliminary tests using simplified scenario).
- Direct plotting of data cells – quadrilaterals only at first, but hexagonal and triangular cells are in the planning stages.
- Continuous zooming from many data cells per pixel to many pixels per data cell.
- Each data cell independently rendered – **parallel support considered in design.**
- Intend to provide visualization capabilities suitable for data of interest to the ParVis project
- Preliminary version by 6.1.0 beta (hopefully)



Progress in Task-parallel diagnostics: Swift and AMWG

- **Swift-based AMWG diagnostics released to community!**
 - Officially part of version 5.3 of AMWG released in Feb, 2012.
 - Used daily at NCAR
 - Installed on Lens, the DAV cluster at OLCF.
 - Initial tests of using Pagoda (parallel) in place of NCO (serial)

- Swift package recent developments
 - Working on parallel invocation of multiple MPI and OpenMP applications (for using Pagoda in AMWG)

 - Solved problems with scheduling on Eureka (the DAV cluster at ALCF). Improved scheduling on BG/P and Cray.





Swift version of OMWG diagnostics

- Current OMWG diagnostics use non-free software.
- While building Swift version, convert to OMWG diags to all-NCL
 - 87 scripts converted from IDL to NCL
 - All three top-level OMWG control scripts modified to run NCL-based scripts exclusively
 - Graphics appear very similar with identical color table and level spacing to IDL graphics
 - Annotations are at least as original
 - Numerical output matches the original in all important respects
 - 2 routines written as Fortran shared objects for speed
 - Final testing and cleanup by end of April
 - For comparison, see NCL version at:
http://www.ncl.ucar.edu/Applications/popdiag/pd.1981_2005/popdiag.html
vs. original version at
http://www.cesm.ucar.edu/experiments/cesm1.0/diagnostics/b40.20th.track1.1deg.005/ocn_1981-2005-obs/popdiag.html
- Swift version being tested for many different file control options.
 - Already calling NCL scripts and can easily update to final versions.

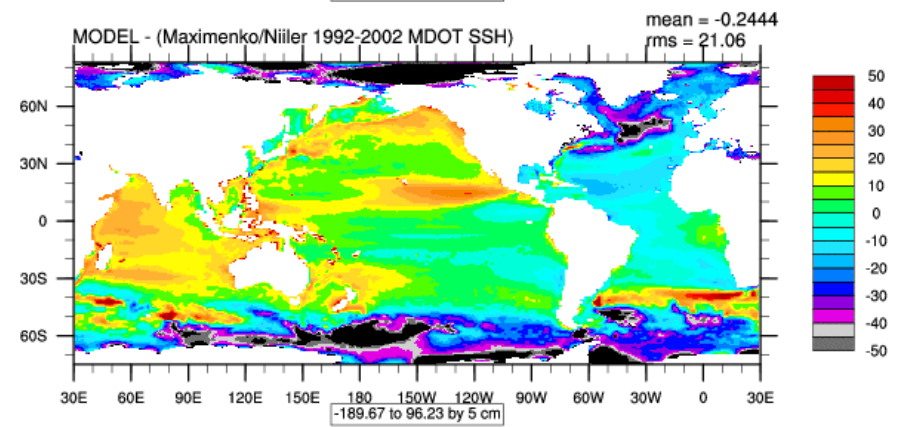
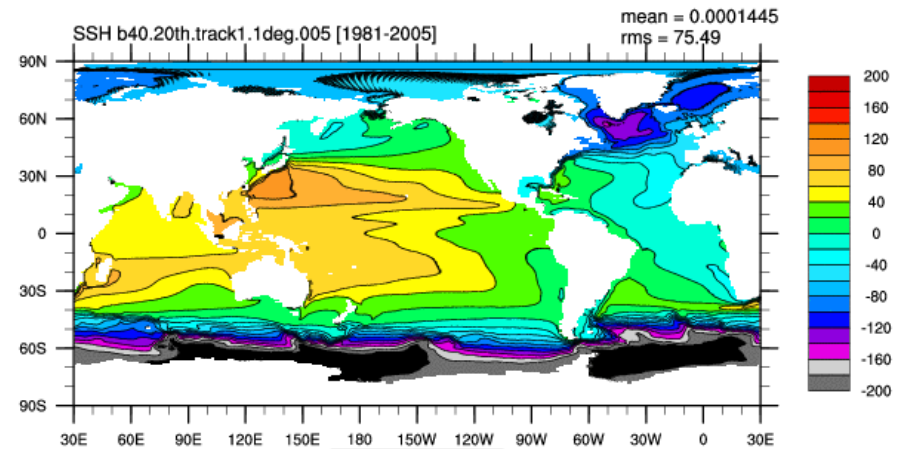


OMWG diagnostics: Sea Surface Height

Original



NCL



Cloud Computing paradigms: MapReduce

- Beginning of FY12: Working “Streaming Hadoop” prototype kernels for averaging
 - Testbed kernel for probability density function (PDF) estimation implemented; applied to problem of time-evolving PDF $f(X,t)$ estimation
 - Publication: “Visualizing climate variability with time-evolving probability density functions, detecting it with information theory,” to appear Workshop on Data Mining in Earth System Science, ICCS 2012
- FutureGrid (cloud computing resource) project granted, awaiting allocation to commence scalability tests and performance studies
- Paper “Mapping Climate Data Analysis onto the Map Reduce Programming Pattern” in preparation.




Using GPUs: Interactive Visualization of Large Geodesic Grid Data

Kwan-Liu Ma, UC Davis


Existing 3D visualization solutions:

- Require a pre-partitioning of each hexagonal cell into multiple tetrahedral cells.
- Do not take advantage of latest GPU features
- Do not offer high-quality rendering



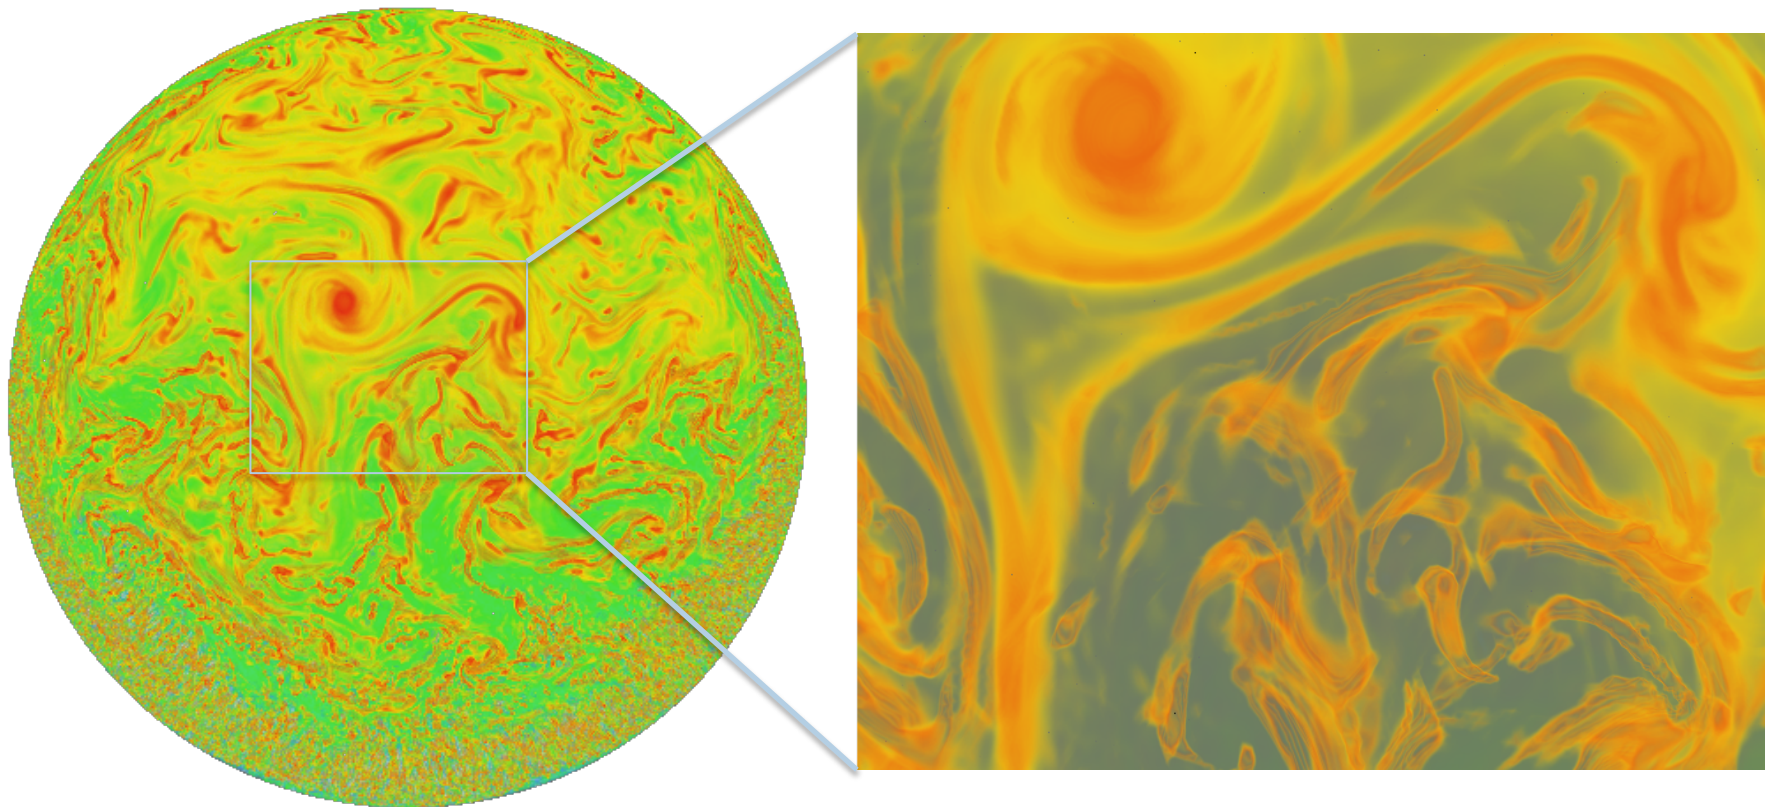
The UC Davis team seeks to provide:

- Advanced visualization of hexagonal grid data
- High quality 3D rendering
- GPU acceleration and parallelization to support Interactive interrogation



Interactive Visualization of Large Geodesic Grid Data

Kwan-Liu Ma, UC Davis



Data: CSU GCRM



parVis


Interactive Visualization of Large Geodesic Grid Data

Kwan-Liu Ma, UC Davis

Further Work:

- Further optimizing performance
- Interactive interrogation
- Rendering on a GPU cluster
- Time-varying data visualization
- Working directly with the GCRM team
- In situ visualization

Data: CSU GCRM



Disk Space: Progress in Compression schemes

- Tested several compression schemes for climate data in first year.
- Need for compression is here now: Long run of 0.10 POP at NCAR has to output less data because can't fit on disk.
- Improve end-to-end performance of working with compressed data.
 - Reduce disk overhead
 - Reduce communication overhead
 - Offset computation overhead of compression
- Keeping compression easy to use
 - Application transparent if possible
 - Minimize software layers to be changed



Compression: File System Based Approach

- Change the file systems to return decompressed data
 - Caching to reduce computation overhead
 - Pipelining and Prefetching to reduce latency
- Advantage
 - Application transparent
- Disadvantage
 - Does not reduce communication cost
 - Need to change file systems



PnetCDF Based Approach

- Fetch compressed data through MPI-IO
- Advantages
 - Reduce disk overhead
 - Reduce communication overhead
- Disadvantage
 - Challenging when PnetCDF accesses and data compression are not aligned
 - Pipelining is difficult
- Implemented a proof of concept prototype and performed some preliminary measurements
 - Read a 2.7 gb netcdf file with uncomopressed data, 39.454 seconds, with compressed data, 27.429 second



Outreach to community

- Convened a session at Fall 2011 AGU with other Viz PI's (Williams, Bethel) "Challenges in Analysis and Visualization of Large Earth Science Data Sets"
 - Oral presentation on ParCAL/ParNCL
 - Poster on Swift AMWG
- Presented ParVis and ParCAL at UCAR Software Engineering Assembly, February, 2012
- Continue to update website and parvis-ann mailing list.



ParVis Near-term plans

- ParCAL/ParNCL
 - Continue implementing more NCL functions
 - Put priority on functions in new NCL analysis routines written by CSSEF Atmosphere Team
 - Read in and operate on CAM-SE unstructured grids.
 - Define unstructured grid functions for common atmospheric analysis such as “zonal average”
- Task-parallel diags
 - Finish OWMG diags and release in time for CESM Tutorial
 - Begin parallelization of LMWG diags to support CSSEF land modeling
- Outreach:
 - Poster accepted to EGU (April).
 - DOE Computer Graphics Forum (April)
 - Present at CESM Workshop (June) and CESM Tutorial
 - Submit paper to **IEEE Symposium on Large-Scale Data Analysis and Visualization** (Oct, 2012)





parVis

<http://trac.mcs.anl.gov/projects/parvis>



parVis