Evaluating Performance and Scalability of Candidate Dynamical Cores for the Next Generation Global Prediction System

John Michalakes [1], Rusty Benson [2], Tom Black [1], Michael Duda [3], Mark Govett [4], T. Henderson [4], P. Madden [4], George Mozdzynski [5],Alex Reinecke [6], William Skamarock [3], Ratko Vasic [1]

[1] NOAA National Centers for Environmental Prediction
[2] NOAA Geophysical Fluid Dynamics Laboratory
[3] National Center for Atmospheric Research
[4] NOAA Earth System Research Laboratory
[5] European Centre for Medium-Range Weather Forecasts
[6] Naval Research Laboratory

2015 NCAR Multi-core 5 Workshop
September 16-17

NOAA/NWS/Environmental Modeling Center

Outline

- NGGPS Program
- Dycore Test Plan
- Advance Computing Evaluation Benchmarks
- Summary

Next Generation Global Prediction System

- NGGPS is a program within National Weather Service 5 yr R2O Effort
- Design, develop, implement in operations a fully coupled atmos/ocean/wave/land/aerosol global prediction system in 2020 (PM: Fred Toepfer)
  - Explicit simulation of moist convention on global scale
  - New non-hydrostatic Dynamical Core to replace current Global Spectral/Semi-Lagrangian hydrostatic dynamical core in GFS
  - Take advantage of advances in High Performance Computing
- http://www.nws.noaa.gov/ost/nggps/

NGGPS Test Plan Timeline

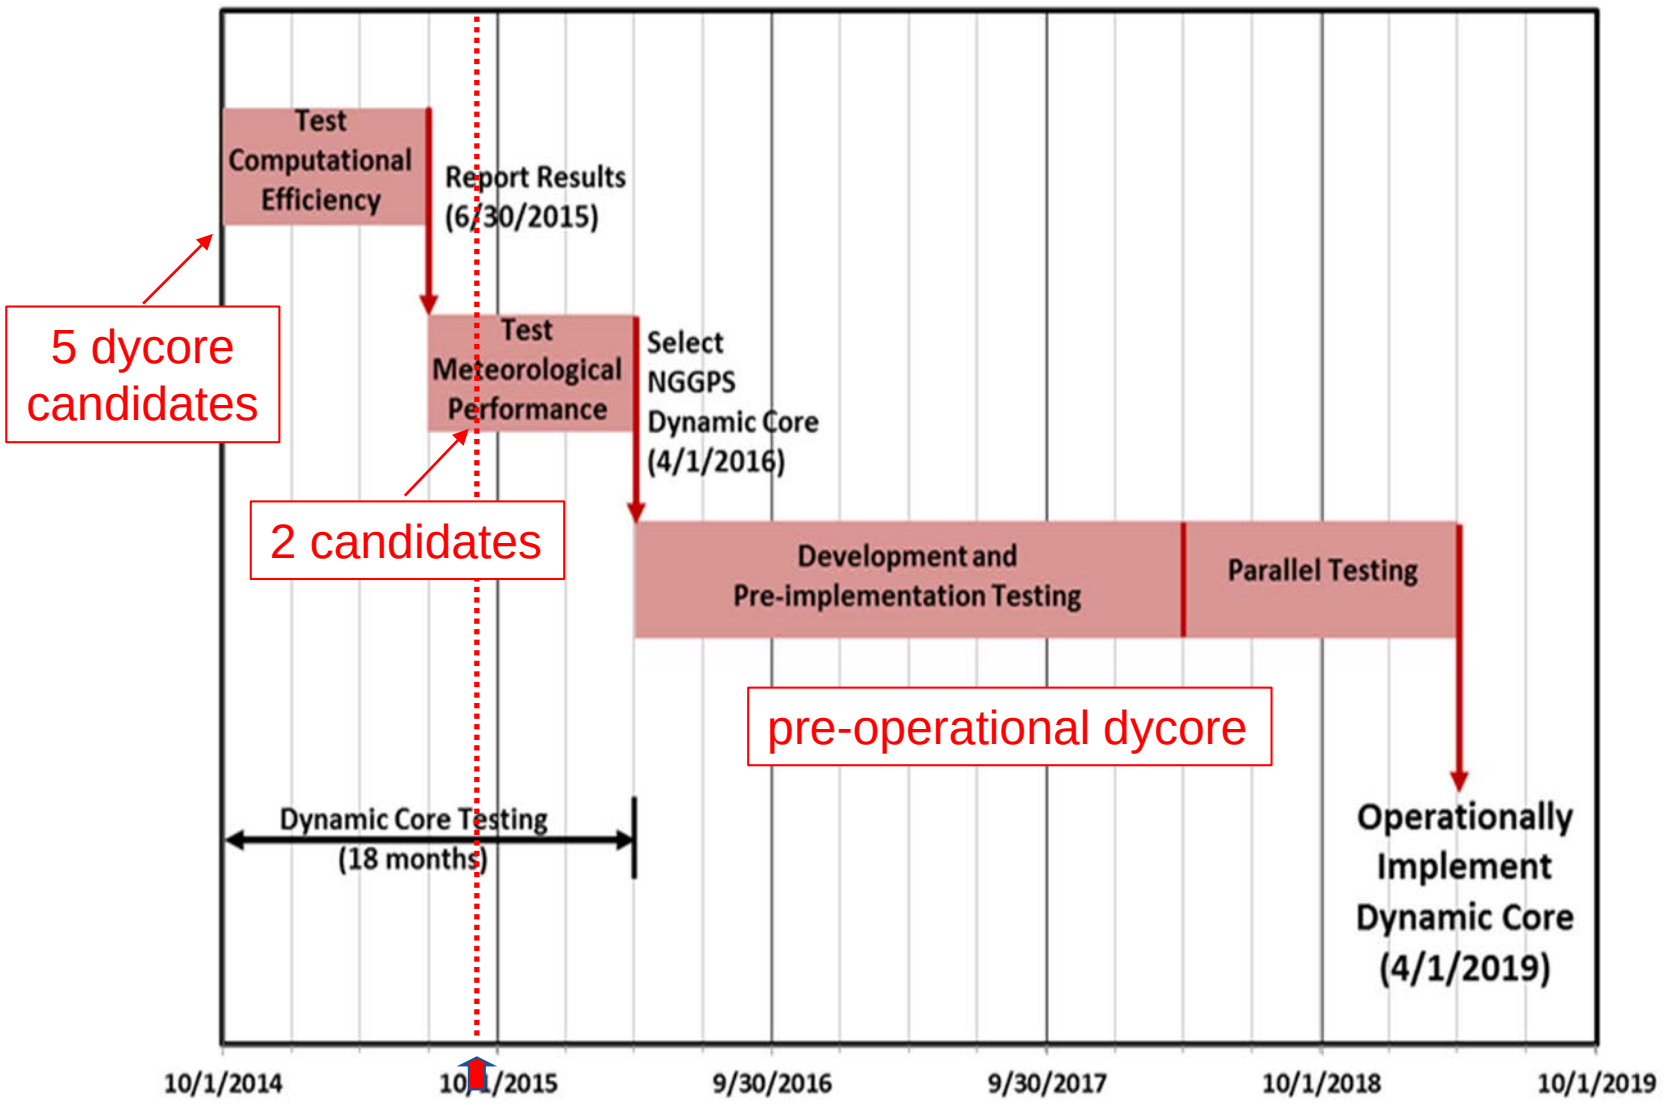

Advanced Computing Evaluation Committee

- AVEC formed August 2014 to evaluate and report on performance, scalability and software readiness of five NGGPS candidate dycores:

Model	Organization	Numeric Method	Grid
NIM	NOAA/ESRL	Finite Volume	Icosahedral
MPAS	NCAR/LANL	Finite Volume	Icosahedral/Unstructured
NEPTUNE	Navy/NRL	Spectral Element	Cubed-Sphere with AMR
HIRAM/FV-3	NOAA/GFDL	Finite Volume	Cubed-Sphere, nested
NMM-UJ ***	NOAA/EMC	Finite difference	Cubed-Sphere
~~GFS-NH *~~	~~NOAA/EMC~~	~~Semi-Lagrangian/Spectral~~	~~Reduced Gaussian~~
IFS (RAPS13)**	ECMWF	Semi-Lagrangian/Spectral	Reduced Gaussian

* Current operational baseline, non-hydrostatic option under development, No version of GFS was available for AVEC tests

** Guest dycore, hydrostatic, GFS proxy

*** NMMB replaced by NMM-UJ

Workloads

- Baroclinic wave case from HIWPP non-hydrostatic dycore testing (DCMIP 4.1)
  - Added 10 artificial 3D tracer fields to simulate cost of advection
  - Initialized to checkerboard pattern to trigger cost of monotonic limiters
- 13 km workload
  - Represent current and near-term global NWP domains
  - Measure performance of the code with respect to operational time-to-solution requirement (8.5 minutes/forecast day)
- 3 km workload
  - Represent workloads that might be in operations within lifetime of NGGPS
  - Measure ability to scale to efficiently utilize many times greater computational resources

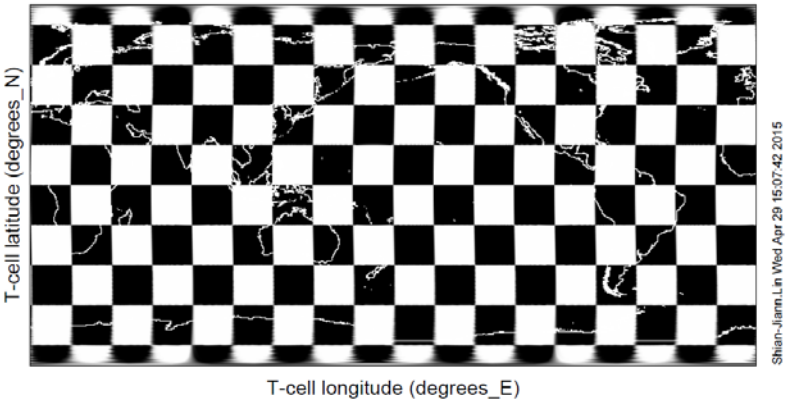

Checkerboard tracer initialization pattern after one hour FV3 integration. Image provided by S. J. Lin, NOAA/GFDL

Benchmark Configurations

		NH-GFS (Baseline) *	FV3	MPAS	NIM	NMMB-UJ	NEPTUNE	IFS (RAPS13) *
Nominally 13km	Resolution	13 km (TL1534)	~12 km (C768)*	12km *	13.4 *	13 km	12.71 km *	12.5 km (Tc799)
	Grid Points	3072x1536 (unreduced) 3,126,128 (reduced)	6x768x768 3,538,944	4,096,002 **	3,317,762	6x768x768 3,538,944 *	3,110,402 **	3,336,946 (reduced)
	Vertical Layers *	128	127 **	127 ***	128	128	127 ***	137
	Time Step	TBD	600s (slow phys) 150s (vertical, fast phys) 150/10 (horiz. acoustic)	72 s (RK3 dynamics) 12 s (acoustic) 72 s (RK3 scalar transport)	72 s	24 s **	75 s (advective), 15 s (sound) ****	450
Nominally 3km	Resolution	3 km (TL6718)	~3 km (C3072) *	3km	3.3 km **	3 km	3.13 km *	3.125 km (Tc3199)
	Grid Points	13440x6720 (unred.) 59,609,088 (reduced) **	6x3072x3072 56,623,104	65,536,002	53,084,162	6x3072x3072 56,623,104 *	61,440,000 **	51,572,436 (reduced)
	Vertical Layers *	128	127 **	127 ***	128	128	128	137
	Time Step	TBD	150 s (slow phys) 37.5 s (vertical, fast phys) 37.5/10 s (horiz. acoustic)	18 s (RK3 dynamics) 3 s (acoustic) 18 s (RK3 scalar transport)	18 s	6 s **	15 s (slow RK3 dyn.) 2.5 s (fast dyn.)	120

Table A3-1. Model-specific Benchmark Configurations

Computational Resources

Note: Conventional Multi-core Systems only

- Edison: National Energy Research Scientific Computing Center (DOE/NERSC)
  - 4M core hours in two sessions totaling 12 hours of dedicated machine access
  - 133,824 cores in 5,576 dual Intel Xeon Ivy Bridge nodes (24 cores per node)
  - Cray Aries with Dragonfly network topology
  - https://www.nersc.gov/users/computational-systems/edison/configuration
- Stampede: Texas Advanced Computing Center
- Pleiades: NASA Ames Research Center

Caveats

- The performance and scaling results in this report are a snapshot in time of NWP software that is under active development. The test workloads are based on an idealized atmospheric case that does not include physics.
- The choice of time step for the idealized benchmark runs was best-guess of what would be needed for full-physics real-data forecasts on the part of the modeling groups. In adjusting benchmarking results to the operational speed requirement, we also assumed that dynamics represents half the run time of a full-physics model.
- Benchmarks were compute-only**.** AVEC did not evaluate important aspects of performance such as I/O, initialization costs, or other factors that would not represent full physics realizations of the models. Such testing will occur in future Phase-2 evaluations under the NGGPS test plan.
- AVEC evaluated model performance with no consideration for solution quality. Each candidate model's benchmarks were conducted with the same formulation and configuration used to run the idealized test cases under the High Impact Weather Prediction Program (HIWPP) non-hydrostatic dycore evaluations (led by Jeff Whitaker, NOAA/ESRL).

AVEC Level-1 Evaluations: Performance

- Performance:
  - Number of processor cores needed to meet operational speed requirement with 13-km workload
  - Candidate rankings (fastest to slowest): (1) NMM-UJ, (2) FV3, (3) NIM, (4) MPAS, (5) NEPTUNE

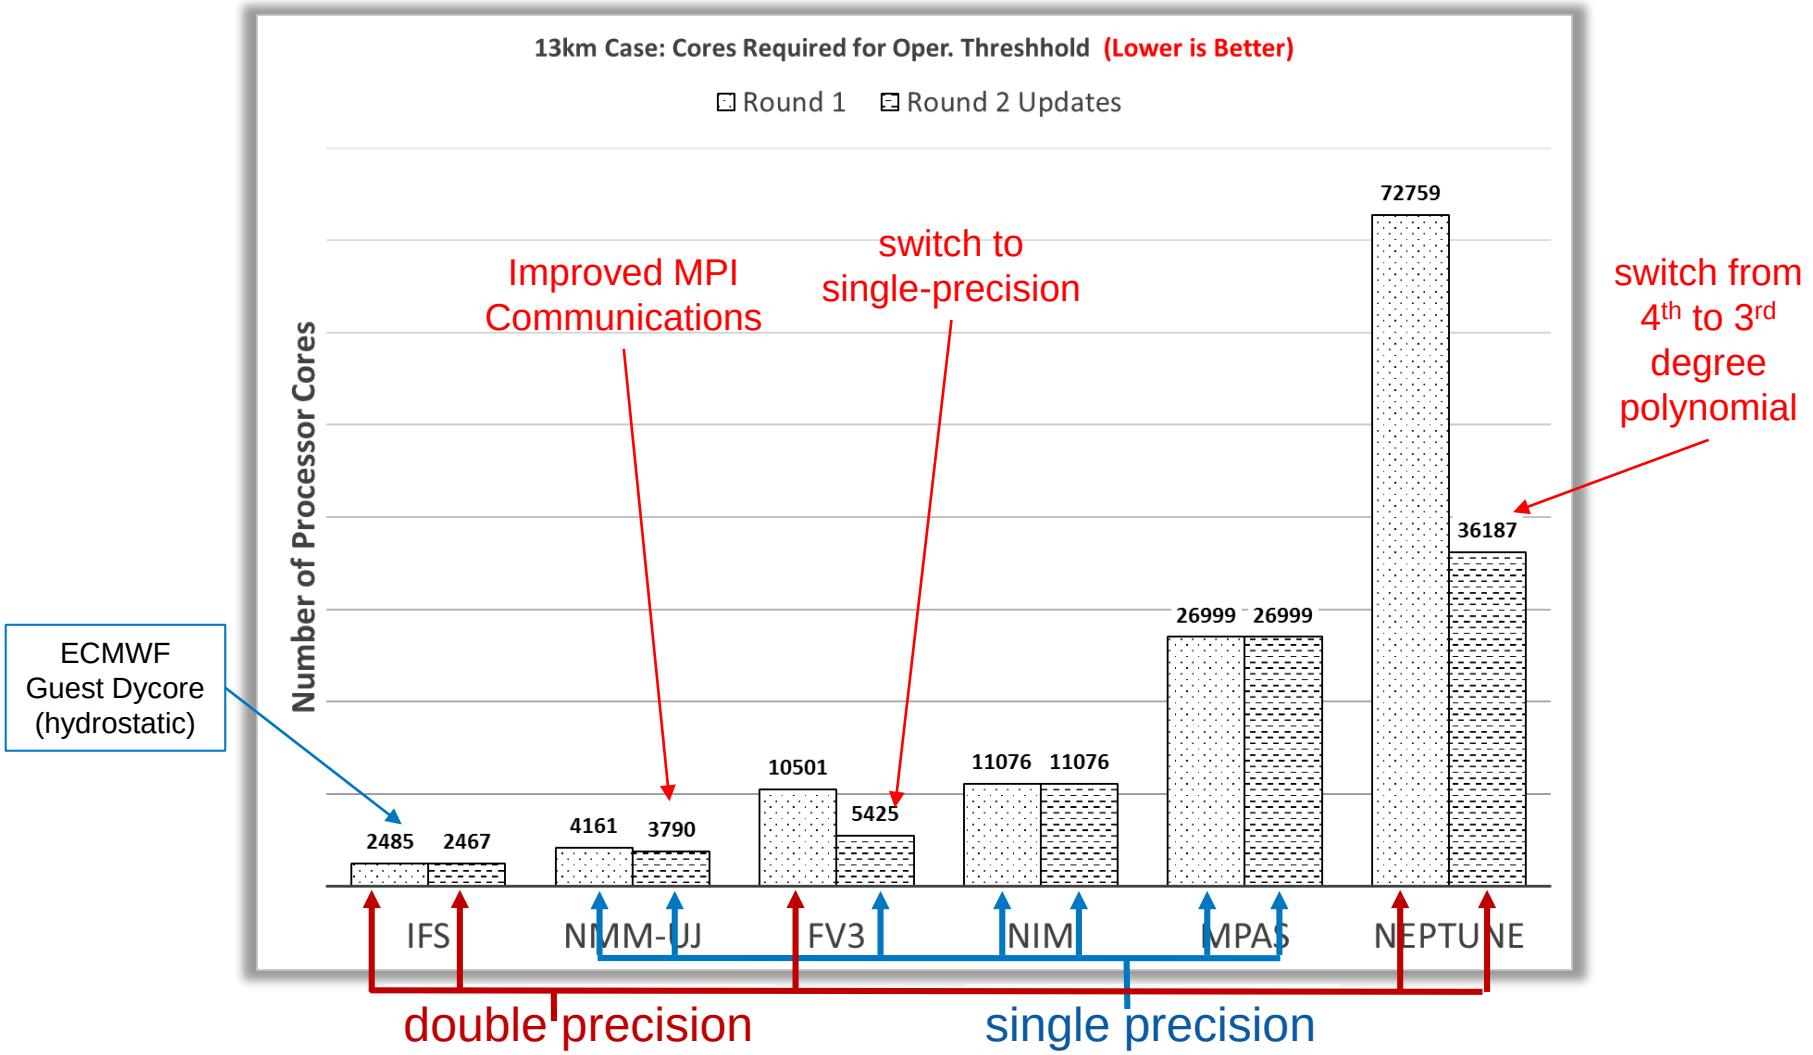

AVEC Level-1 Evaluations: Scalability

- Scalability: ability to efficiently use large numbers of processor cores
  - All codes showed good scaling.
  - Candidate rankings (scalability): (1) NEPTUNE, (2) MPAS, (3) NIM, (4) FV3, (5) NMM-UJ

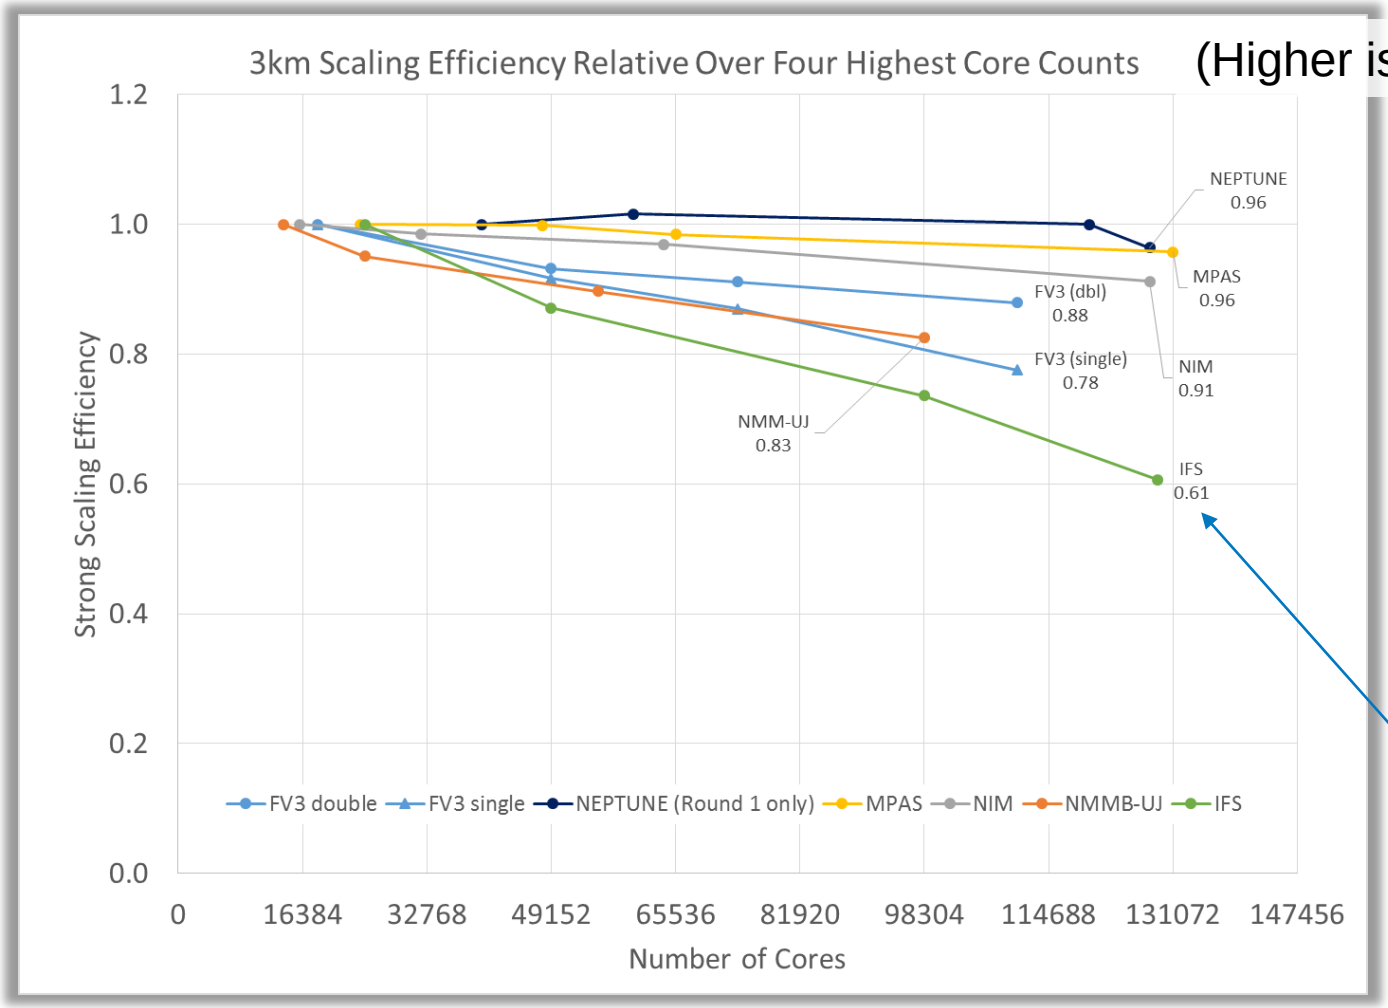

NGGPS Phase 1 Testing
Project Summary Assessment

	Idealized Tests	3-km, 3-day forecasts	Performance	Scalability	Nesting or Mesh Refinement	Software Maturity
FV3						
MPAS						
NIM						
NMM-UJ						
NEPTUNE						

Selected to progress to Phase-II Testing

 Meets or exceeds readiness for needed capability

 Some capability but effort required for readiness

 Capability in planning only or otherwise insufficiently ready

Idealized Testing Summary

- ***FV3, MPAS*** produced highest quality solutions overall
  - More similar to each other than other models for all tests
- ***NIM*** produced reasonable mountain wave and supercell solutions
  - Excessive noise near grid scale in baroclinic wave solution
  - Full physics forecasts excessively damped
- ***NEPTUNE*** was not able to produce full physics 3-km forecasts
  - Baroclinic wave too smooth, 4-km supercell not split by 90 mins
- ***NMM-UJ*** did not produce realistic solutions for the mountain wave and supercell tests
  - Vertical velocity fields from full physics forecasts did not show signatures expected from resolved convection

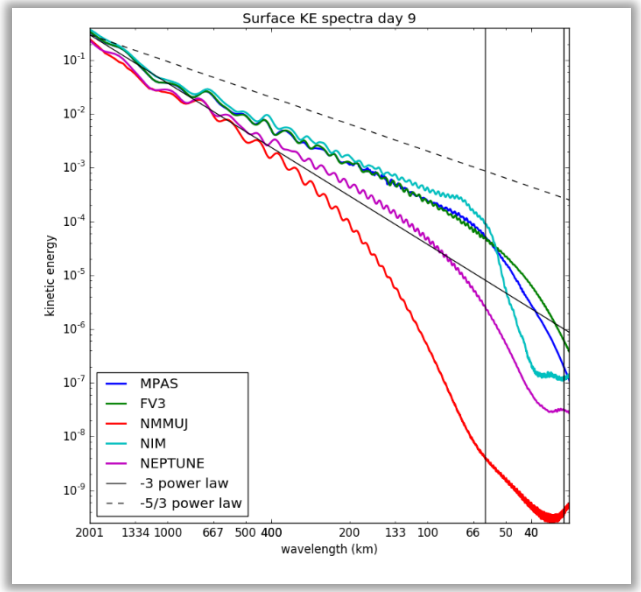

Baroclinic Wave KE Spectrum (surface, day 9, 15-km resolution)

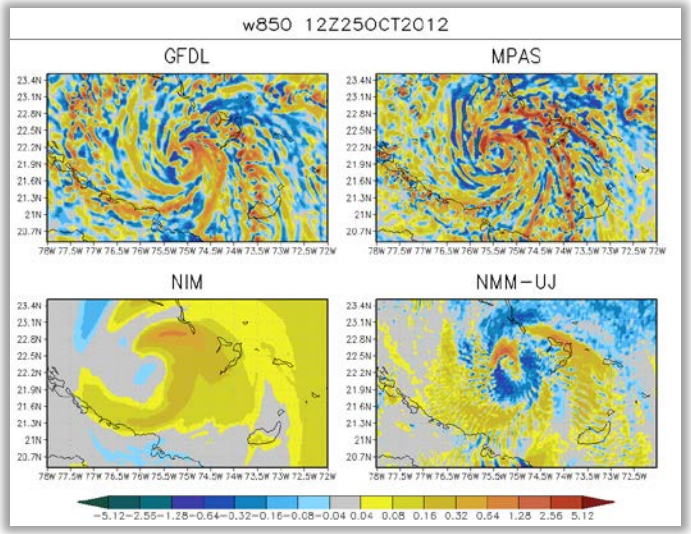

Hurricane Sandy (w at 850 hPa)

Summary

http://www.nws.noaa.gov/ost/nggps

- NGGPS providing opportunity to set technical and scientific course for National Weather Service operations for next 10-15 years
  - Represent the earth system with unprecedented breadth and fidelity
  - Increase participation of research community
  - Take advantage of increasing computing power towards exascale
- Next steps
  - Phase-2 testing underway for MPAS and FV3
    - More detailed meteorological testing
    - Computational testing with performance and approaches to mesh refinement/nesting
  - More detailed evaluation of software readiness for next generation architectures
- Phase-1 NGGPS Testing Reports:
  - AVEC performance and scaling report: “NGGPS Level-1 Benchmarks and Software Evaluation” (4/30/2015)
  - Reports on solution accuracy testing:
    - “HIWPP non-hydrostatic dynamical core tests: Results from idealized test cases.” **Jeffrey Whitaker**, NOAA/ESRL
    - “HIWPP non-hydrostatic dynamical core tests: Results from 3-km, three-day simulations.” **Jeffrey Whitaker and Philip Pegion** NOAA/ESRL

Acknowledgements

http://www.nws.noaa.gov/ost/nggps

- Dycore solution quality testing
  - NOAA/OAR's High Impact Weather Prediction Program (HIWPP)
  - Jeff Whitaker, NOAA/ESRL
- Benchmark Computing Resources
  - National Energy Research Scientific Computing Center (DOE/NERSC)
  - Texas Advance Computing Center (NSF/TACC)
  - NASA Advanced Supercomputing Division
- IFS model and support for AVEC benchmarks
  - European Centre for Medium-Range Weather Forecasts

- References

  - Donner L. J, and coauthors, 2011: The Dynamical Core, Physical Parameterizations, and Basic Simulation Characteristics of the Atmospheric Component AM3 of the GFDL Global Coupled Model CM3. *J. Climate*, **24**, 3484–3519. doi: http://dx.doi.org/10.1175/2011JCLI3955.1
  - Giraldo, F., Kelly, J. , and E. Constantinescu, E. M., 2013: Implicit-explicit formulations of a three-dimensional non-hydrostatic unified model of the atmosphere (NUMA). *SIAM Journal on Scientific Computing*, **35(5)**, B1162-B1194. doi: 10.1137/120876034
  - Janjic, Z., and R.L. Gall, 2012: Scientific documentation of the NCEP non-hydrostatic multistage model on the B grid (NMMB). Part 1 Dynamics. NCAR Technical Note NCAR/TN-489+STR, doi: 10.5065/D6WH2MZX.
  - Lin, S.-J., 2004: A vertically Lagrangian Finite-Volume Dynamical Core for Global Models. *Mon. Wea. Rev*., **132**, 2293-2307. doi: http://dx.doi.org/10.1175/1520-0493(2004)132<2293:AVLFDC>2.0.CO;2
  - Skamarock, W, M. Duda, L. Fowler, S.-.H Park and T. Ringler, 2012: A Multi-scale Nonhydrostatic Atmospheric Model Using Centroidal Voronoi Tesselations and C-Grid Staggering. *Mon. Wea. Rev.*, **240**, 3090-3105, doi: http://dx.doi.org/10.1175/MWR-D-11-00215.1