#4: Performance Benchmark: Methodology

- GFS physics runs with double (64b) floating point precision
- Configurations same as for retro forecasts
- 3 nominal resolutions: 15 km, 13 km, 11 km; 63 levels (so differences in effective resolution could be accounted for). Benchmark parameters agreed to by NCAR and GFDL
- Dedicated access to Cori system at NERSC (similar to Luna/Surge); runs conducted on otherwise empty machine
- Metric: Number of processors required to achieve 8.5 minutes per day simulation rate
- Multiple runs varying numbers of processors to straddle 8.5 min/day simulation rate
- Also tested were:
  - Efficiency of mesh refinement strategies (using configuration for criteria #5)
  - Performance with 15 and 30 extra tracers

#4 Performance Benchmark: KE Spectra (Effective Resolution)

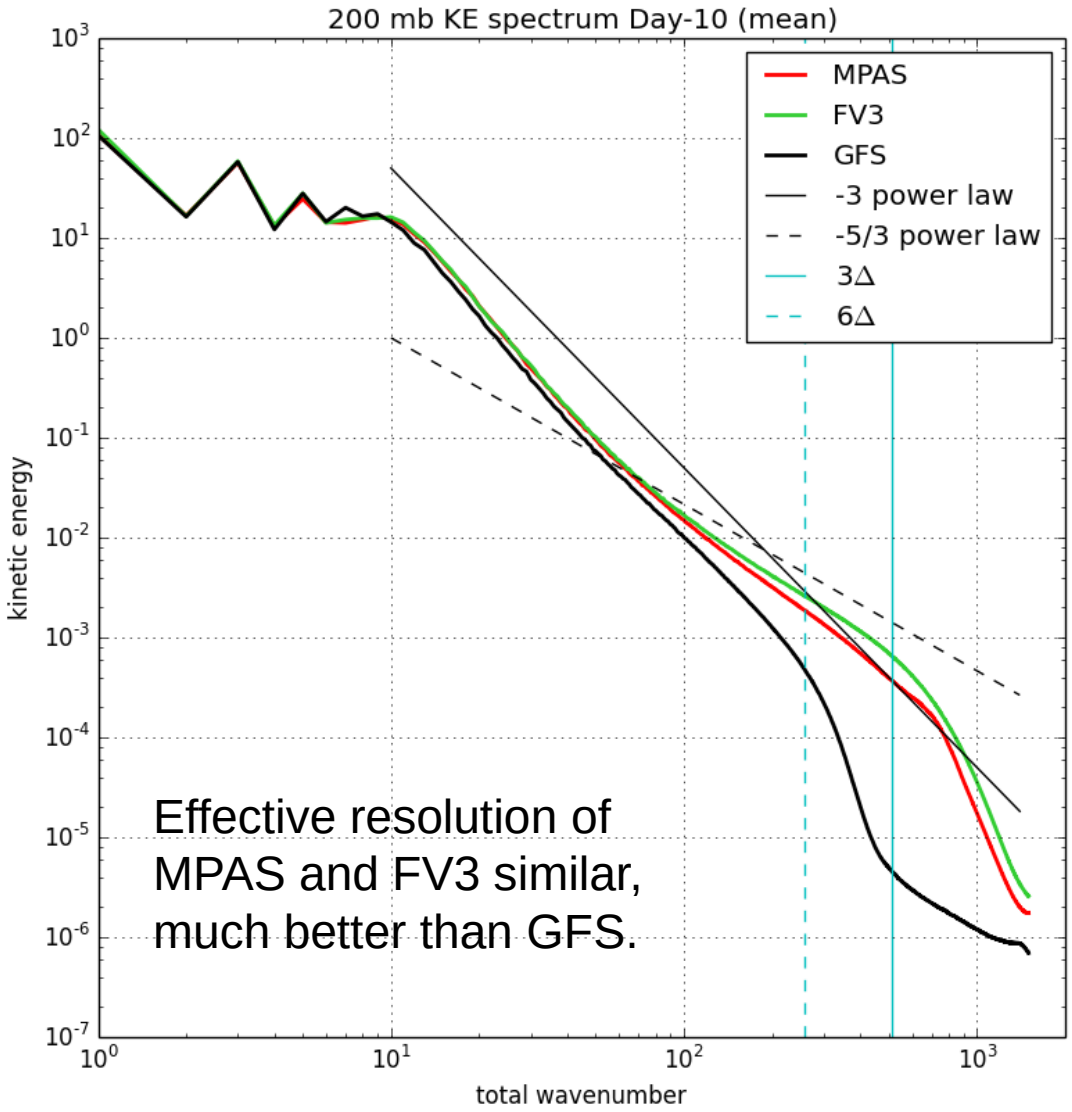

#4: Performance Benchmark Results (J. Michalakes)

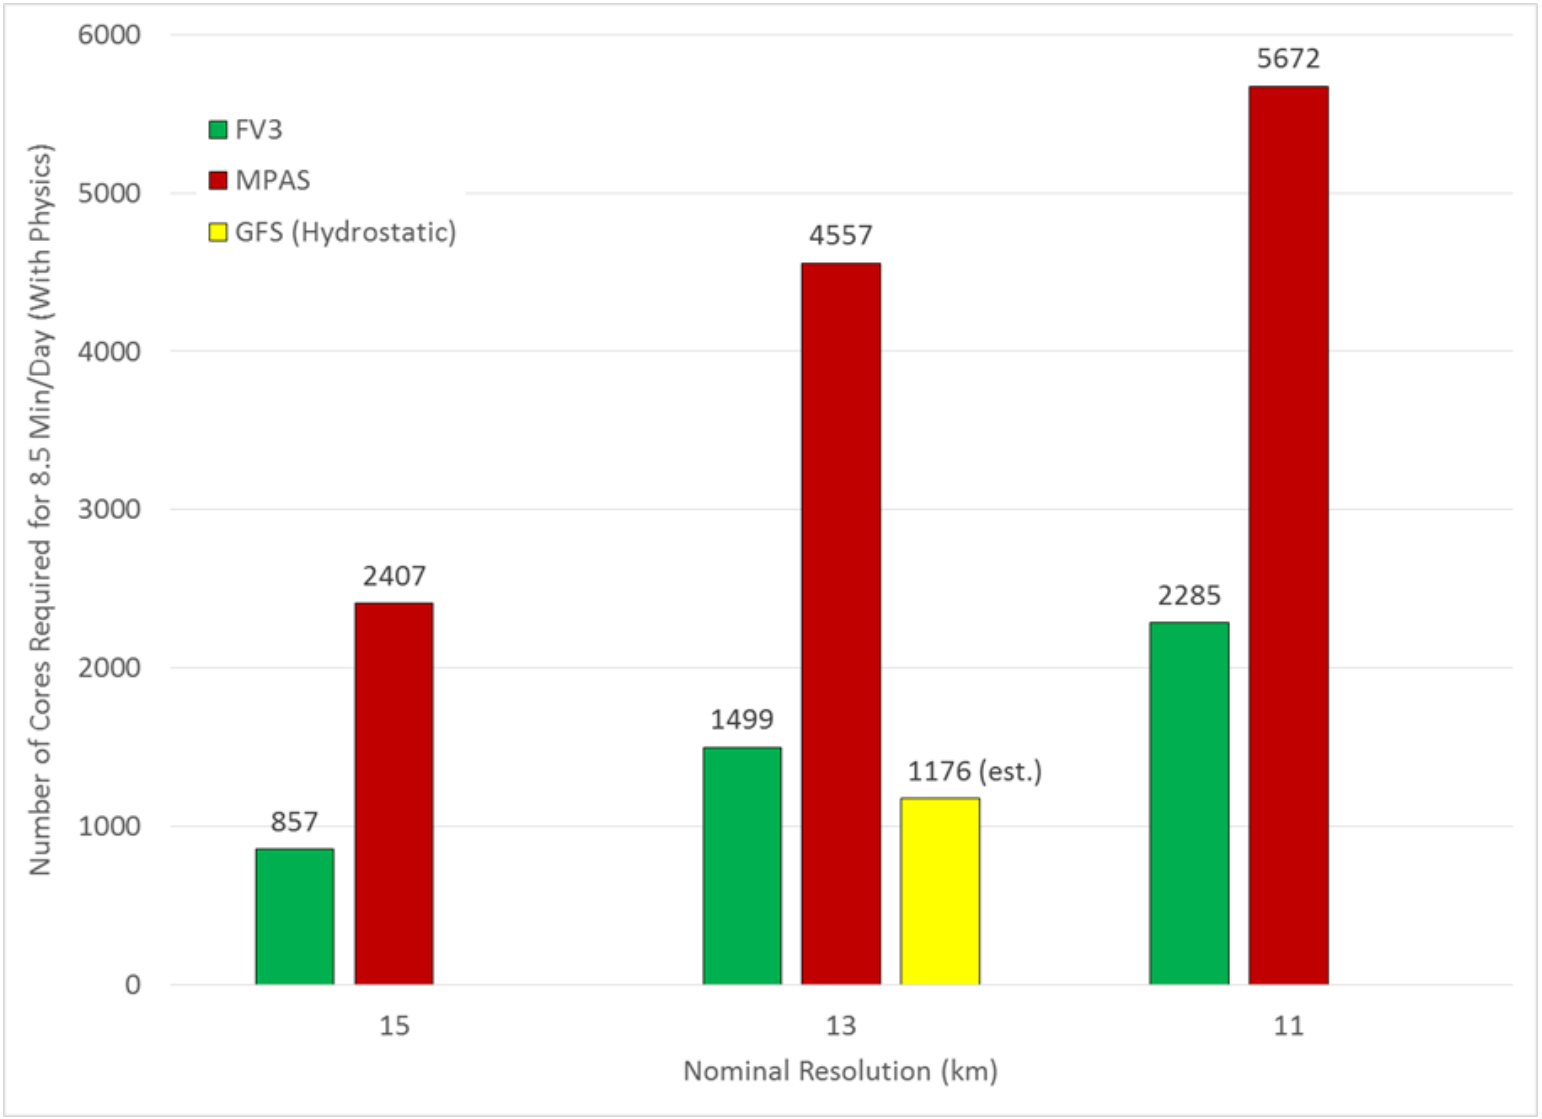

#4: Performance Benchmark Results: Estimated Spectral Slope

Vertical lines are 4dx and 8dx.

Horizontal lines are -3 and -5/3.

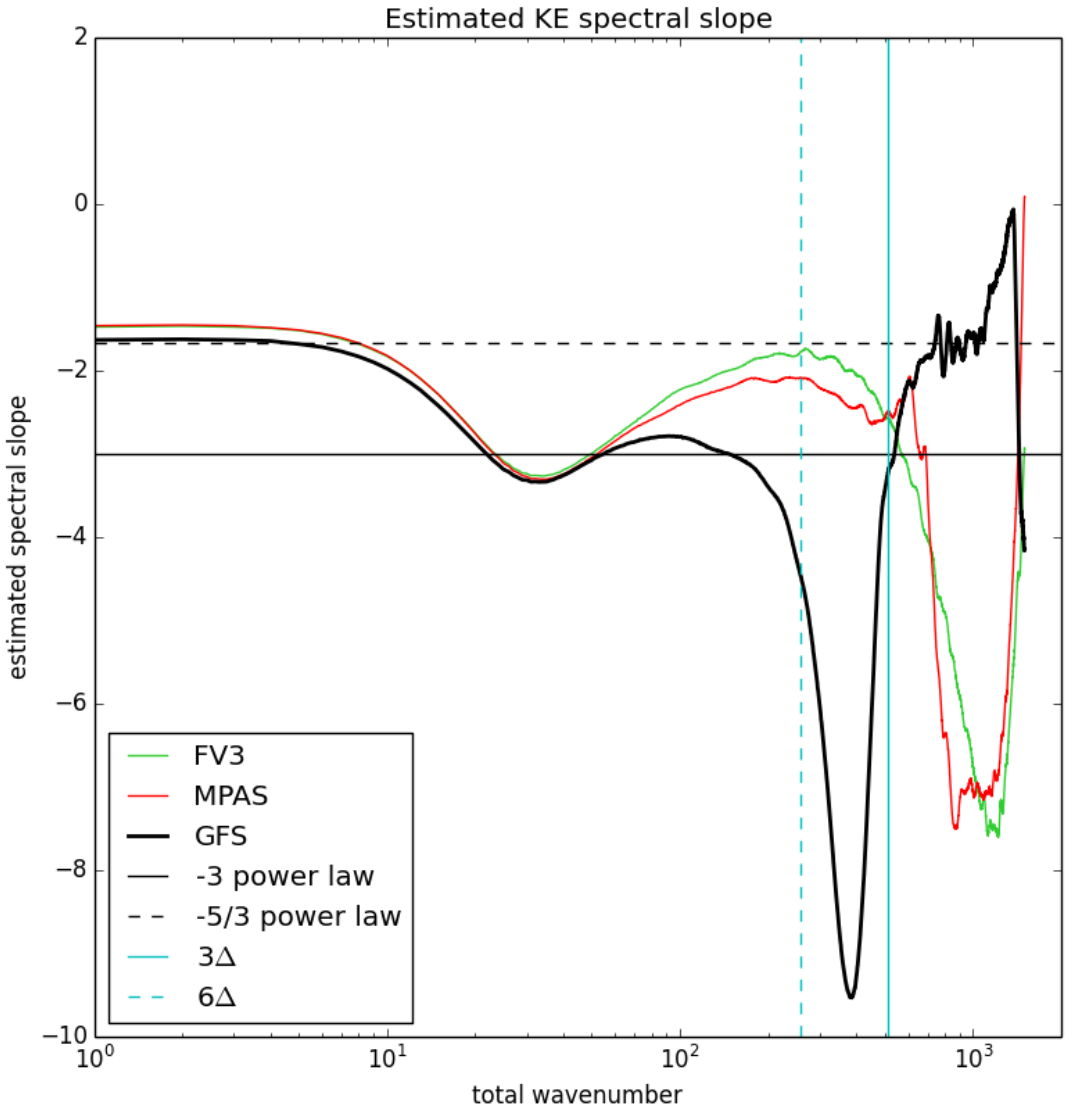

#4: Performance Benchmark Results: Configurations

Eval. Criterion #4 -- Performance with GFS Physics		
	FV-3	MPAS
Nominal resolution (km)	13.03 (equat.), 12.05 (avg.)	13
Grid Points	3,538,944	3,504,642
Vertical Layers	63	63
Time Step (sim. sec)	112.5 (dyn.), 18.75 (acous.)	75 (transport), 37.5 (dynamics), 18.75 (acoustic)
Radiation Time Step	3600	3600
Physics (other) Time Step	225	225
Tracers	3	3
Coarser than nominal resolution (km)	15.64 (equat.), 14.46 (avg.)	15
Grid Points	2,547,600	2,621,442
Vertical Layers	63	63
Time Step	225 (dyn.), 22.5 (acous.)	90 (transport), 45 (dynamics), 22.5 (acoustic)
Radiation Time Step	3600	3600
Physics Time Step	225	180
Finer than nominal resolution (km)	11.72 (equat.), 10.34 (avg.)	11
Grid Points	4,816,896	4,858,092
Vertical Layers	63	63
Time Step	112.5 (dyn.), 16.07 (acous.)	60 (transport), 30 (dynamics), 15 (acoustic)
Radiation Time Step	3600	3600
Physics Time Step	225	180

#4: Performance Benchmark Results: Tracer advection performance

- Tracer advection benchmarks on Cori
  - Measure cost as a function of number of 3D tracer fields
    - Workloads and configuration:
      - 13 km case on number of cores needed for 8-8.5 min/day
    - Baseline: 3 tracer fields
    - Add 15 and 30 artificial tracers
  - Result: cost for full tracer load increased by factor of 2.5 for MPAS versus 1.53 for FV3 compared to baseline.

	Cores	Number of tracers / Minutes			Factor (lowest to highest)
MPAS	4800	3 / 8	18 / 14.6	33 / 19.8	2.5
FV3	1536	3 / 8.14	15 / 9.8	30 / 12.0	1.5 (1.53 adjusted)

Adjustment for FV3 workloads using 15 and 30 tracers **total** *instead of 15 and 30* **additional** *tracers per Test Plan.*

#4: Performance Benchmark Results: Refinement Configuration

Histograms of grid cell size

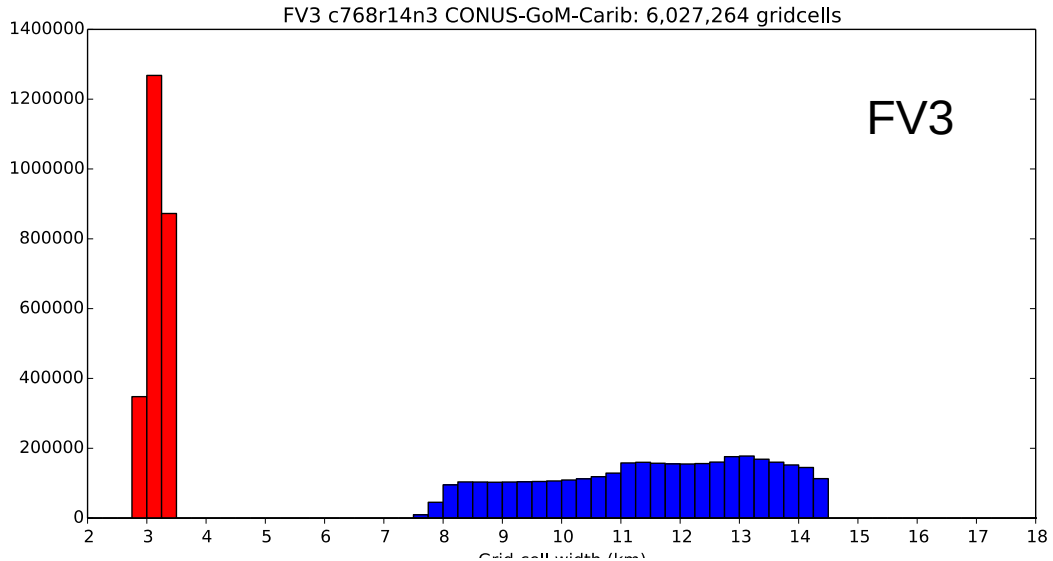

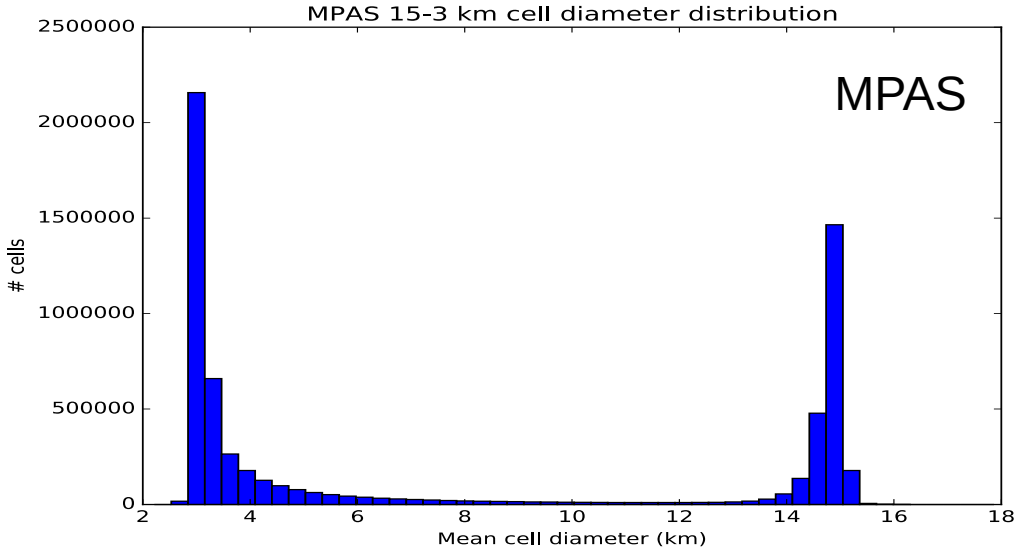

#4: Performance Benchmark Results: Refinement Efficiency

- Part of Criterion #5 evaluation
- How efficient is non-uniform at saving cost compared with uniform 3 km resolution on same number of processors?
- Benchmark and adjust for differences in resolution and area of refinement
- FV3's nesting scheme was more efficient than MPAS's in-place mesh refinement

Definition of nesting efficiency E:

a_g = area of domain (5.101e14 m^2)

a_h = area of refinement (FV3: 2.52e13 m^2 ; MPAS: 2.82e13 m^2)

$r = a_h / a_g$ fraction of domain at high resolution (for uniform res. Domain, r = 1)

dx_L = lowest resolution

dx_H = highest resolution

$C = r\,(dx_L / dx_H)^3\, C_{cellstep} + (1\text{-}r)\, C_{cellstep}$ *(C is "cost")*

$$S_{ideal} = \begin{cases} (dx_L / dx_H)^3 & \leftarrow C_{uniform} \\ r\,(dx_L / dx_H)^3 + 1 - r & \leftarrow C_{refined} \end{cases}$$

(Note: $C_{cellstep}$ factors out)

$$S_{measured} = \frac{T_{uniform}}{T_{refined}}$$

$$E = S_{measured} / S_{ideal}$$

#4: Performance Benchmark Results: Refinement Efficiency (continued)

- Part of Criterion #5 evaluation
- How efficient is non-uniform at saving cost compared with uniform 3 km resolution on same number of processors?
- Benchmark and adjust for differences in resolution and area of refinement
- FV3's nesting scheme was more efficient than MPAS's in-place mesh refinement

	FV3	MPAS
ag (global domain area m^2)	5.101E+14	5.101E+14
ah (high res area m^2)	2.52E+13	2.82E+13
percent of domain in high res r = ah/ag	4.94E-02	5.53E-02
dx low	14	15
dx high	3	3
dx l / dx h	4.67	5.00
(dx l / dx h) ^ 3	101.63	125.00
T-uniform (ideal)	101.63	125.00
T-reduced (ideal)	5.97	7.86
ideal speedup from refinement	17.02	15.91
T_uniform (measured)	345.93	344.65
T_refined (measured)	20.98	34.10
observed speedup from refinement	16.49	10.11
Efficiency	96.9%	63.5%