NCAR

High Resolution Turbulence Resolving Computations---

an example of a very-large simulation

Edward Patton and Peter Sullivan

Franklin Cray XT-4 System

- 9,740 nodes with 19,480 cores
  - 102 Node Cabinets, 16 KWs per cabinet
  - 39.5 TBs Aggregate Memory
- Peak performance: > 100 Tflop/s
- Sustained performance: 19 Tflop/s
- Interconnect: Cray SeaStar2, 3D Torus
  - >6 TB/s Bisection Bandwidth
  - >7 GB/s Link Bandwidth
- Shared Disk: 345 TBs
- Network Connections
  - 4 x 10 Gbps + 16 x 1 Gbps

NCAR LES

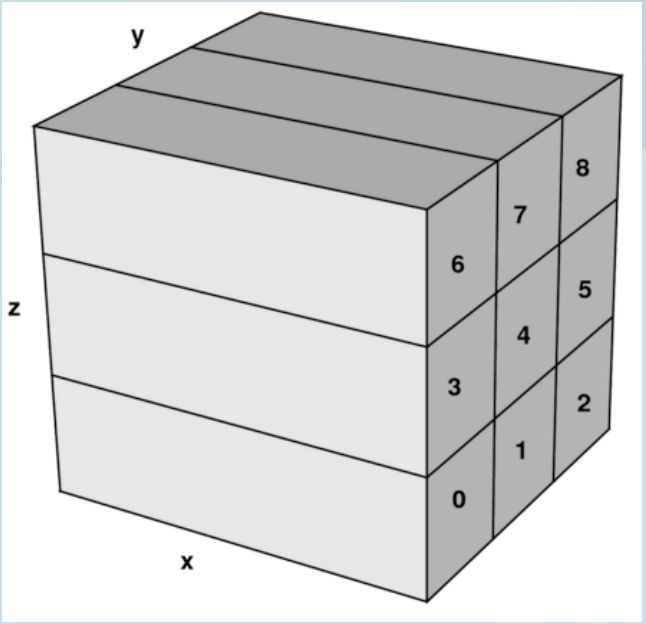

- Pseudospectral in *x* and *y*
- Finite-difference in *z*
- 3rd-order Runge-Kutta in time
- Boussinesq
- Cartesian coordinates
- 2D MPI (broken down in *x* and *z*); each CPU has a cuboid of data. No ALLTOALL communication.

NCAR LES on Franklin

- The largest run: 2048 x 2048 x 2048 grid points
- Domain: 5120 m x 5120 m x 2048 m
- Grid size: 2.5 m x 2.5 m x 1 m
- Run on 8192 processors.
- So far, ran about 500 s simulation time (3000 time steps) which used 36 hr CPU.

NCAR LES on Franklin

- 3D volumes of 6 REAL(KIND=8) variables stored in IEEE binary in a single file using MPI I/O (~ 400 GB per volume).
- Vertical profiles of horizontally-averaged statistics calculated every time step and stored in ‘history’ files.
- 2D instantaneous slices saved for visualizations at specified intervals using MPI I/O.
  - ☛ Possible complications when saving planes of data spanning across the MPI breakdown; every CPU must write only a small piece to the file.

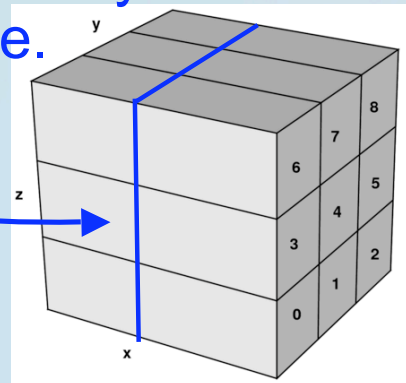

NCAR LES on Franklin

- Key requirements when running on this machine at this resolution
  - Due to file locking issues, I/O on Linux-based machines requires data stored either in:
    - separate files per CPU (requires manipulation after the run)
    - or something like MPI I/O, HDF5, or Parallel-NetCDF to store single files (more work up front, but easier later)
- Integer math used in calculating file pointers in the MPI I/O needs to be (KIND=8)

For our deep-shallow study:

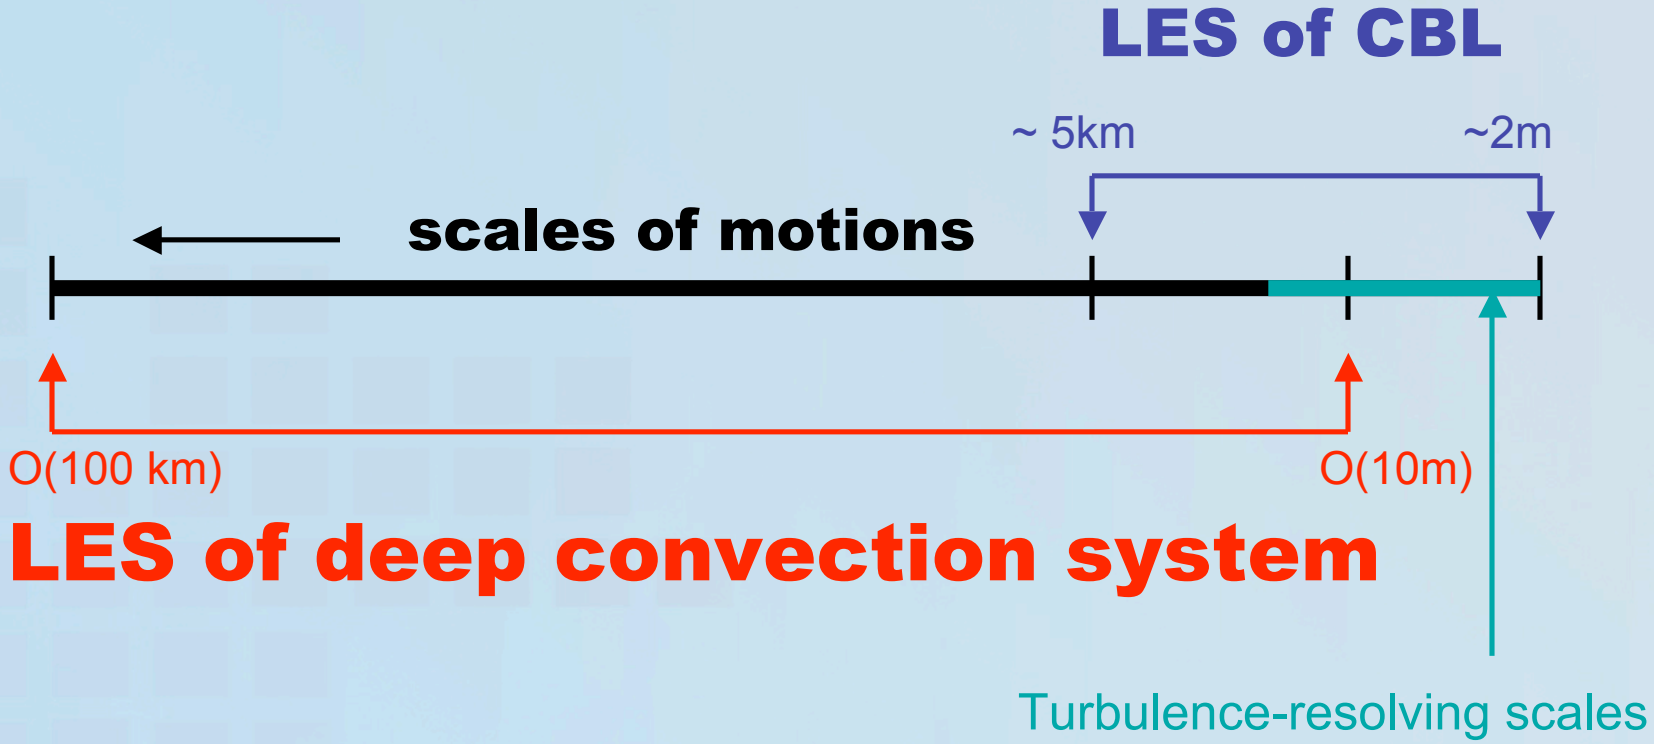

Proposed deep-shallow-convection LES:

- Domain: 400 km x 400 km x 20 km
- Grid size: 100 m x 100 m in x-y; varying in z
- Grid points: 4000 x 4000 x 500 (~2048^3)
- 8 variables (> 6 as the example)
- Which machine(s)? Which model(s)?
- Analyses: statistics? 2D slice data?

NCAR

Rename our theme?

Possible names?

1. Deep and shallow conv, and turb
2. Turbulence in deep conv system
3. Small-scale processes in deep conv system
4. Subgrid-scale processes in cloud-resolving models