

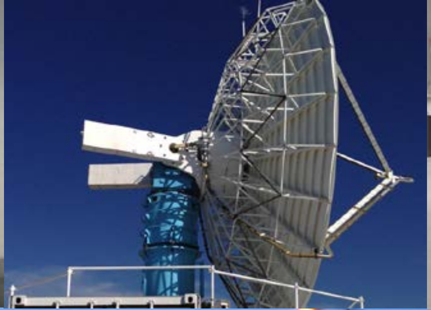


WRF modeling: HPC experiences and cloud challenges

Brian Jewett
University of Illinois
Atmospheric Sciences

WRF Research @ U. Illinois

- Study focus: high impact weather
  - MCSs, severe thunderstorms and tornadoes
  - winter snowstorms • hurricanes
- Resources
  - Teragrid > XSEDE (NCSA, SDSC, PSC, NICS, TACC)
  - *all 'big iron' use, shifting sites; interest in cloud applications*
- Collaborators *UI Atmospheric Sciences*
  - Robert Rauber
    - Professor and head
  - Greg McFarquhar
    - Professor; Director-CIMMS

Perspectives

- All of our use has been on "big iron"
  - very successfully, but with **issues** along the way
- As an atmospheric scientist, am inquiring:
  - how can cloud computing help us w/our **science goals**?
  - can cloud computing reduce the **"80-20" problems**?
- As a XSEDE allocations committee member:
  - what **mix of needs** are best suited for the cloud?
  - can shift to cloud help with high **load** on traditional HPCs?

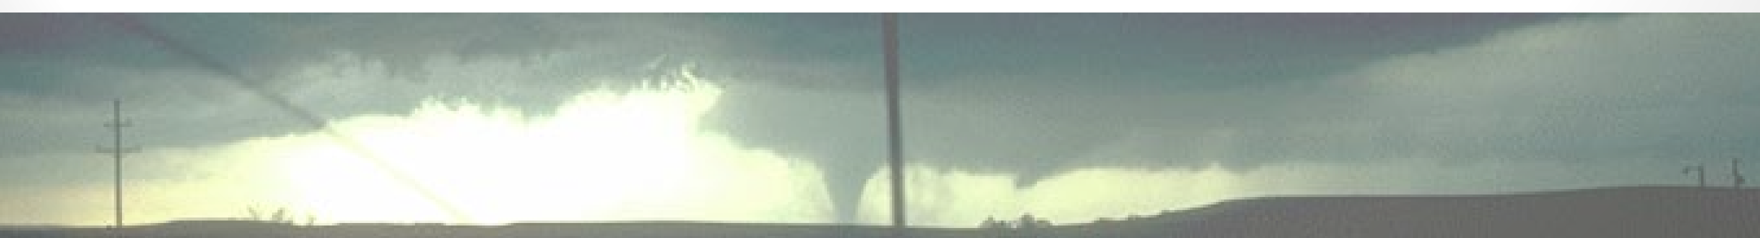

Applications

- Process studies, hypothesis testing
  - Altered physical processes
- High-resolution case studies
  - Trajectory analysis
  - Model vs. observations assessment
  - Place field data in context, e.g. along NCAR/NSF C-130 flight path

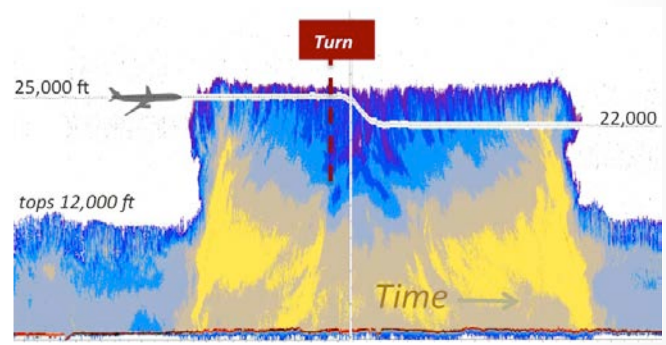

- Recent field programs
  - *PECAN*, Plains Elevated Convection at Night
  - *PLOWS*, ProfiLing Of Winter Storms
  - *SNOWIE*, Seeded & Natural Orographic WIntErtime clouds

PLOWS! Profiling of Winter Storms

Great Lakes blizzard

- *Role of Great Lakes in snow fall depth & distribution*
  - Case: 2011 Chicago Blizzard
  - Grid spacing: 3 km
  - Cores: 512, for 12h
  - I/O: coarse (3h)
  - Cloud applicability: *mixed?*
    - *+ multiple simulations*
    - *+ extended prep, analysis*
    - *- core count + wallclock*

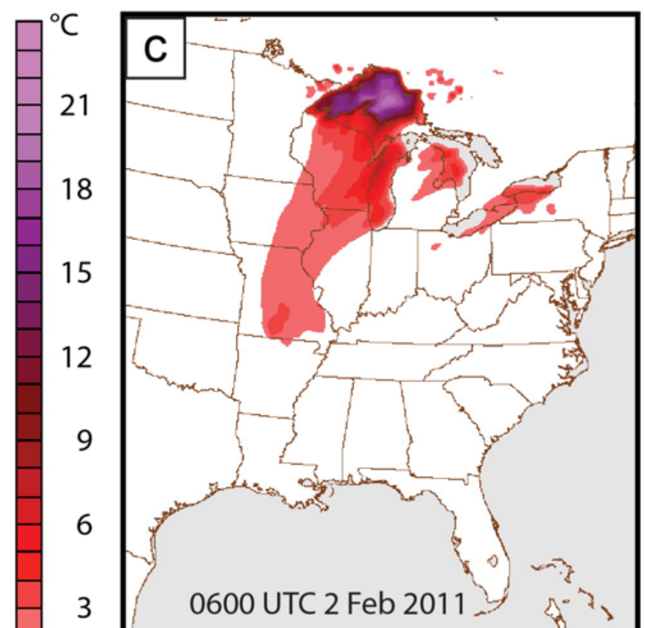

Temp. plume (Lake-none)

PLOWS IOP study

- *Gravity waves and banded precipitation*
  - Case: 12-8-2009 (IOP-10)
  - Domain: ~300x300x300
  - Grid spacing: 3 km x 100m
  - Cores: 512, for 12h
  - I/O: fine (1 min)
  - Cloud applicability: *mixed?*
    - *+ multiple simulations*
    - *- core count + wallclock*

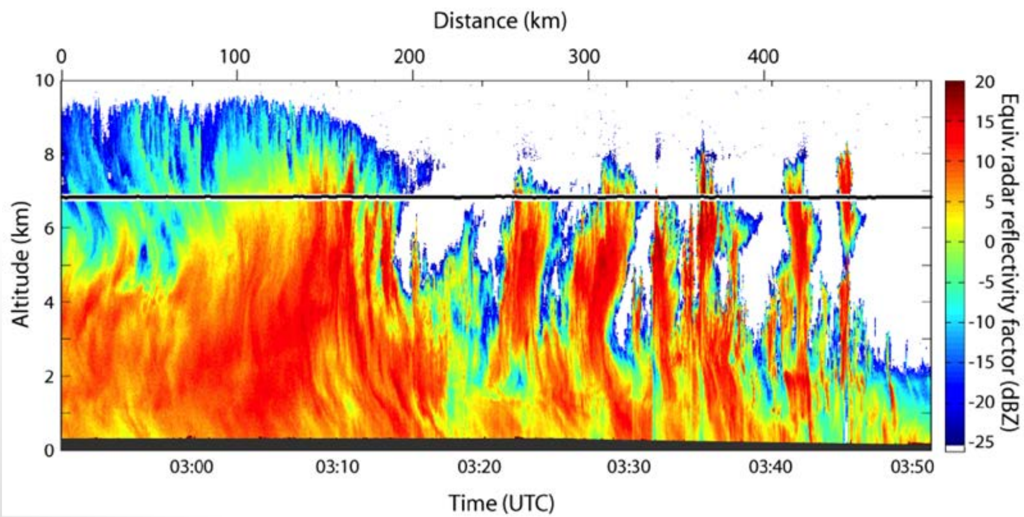

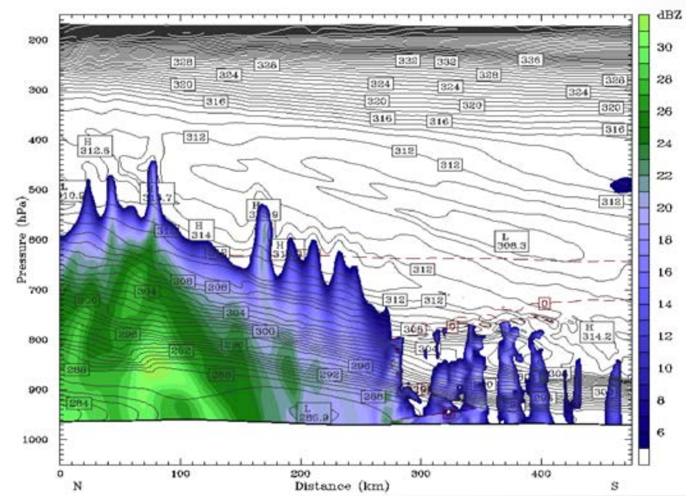

PLOWS! Profiling of Winter Storms

Snow production

- *Cloud-top generating cells*
  - Idealized simulations
  - Parameter sweeps
  - Varied stability, physics
  - Domain: 501x501x199
  - Grid spacing: 100 x 50*m*
  - Cores: 128, for 24h
  - I/O: fine (5 min)
  - Cloud applicability: *low?*
    - core counts, wallclock
    - performance, reproducibility

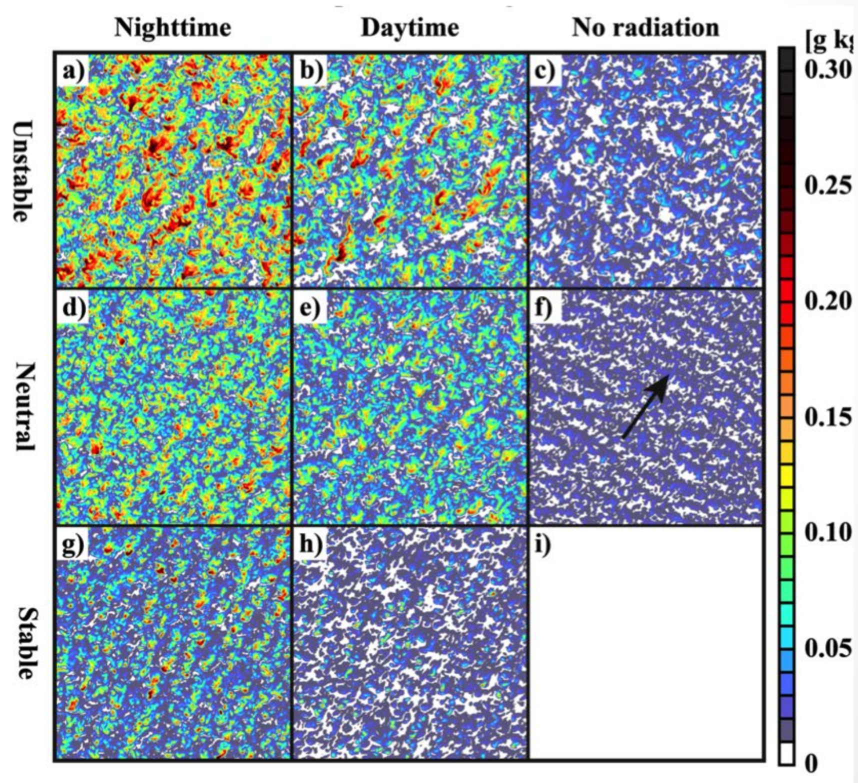

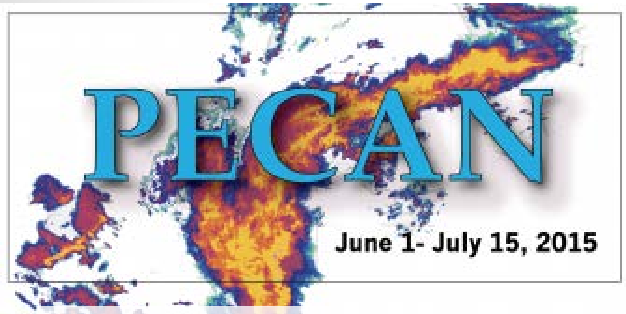

Nocturnal MCSs

- *Real-time simulations in support of field ops*
  - During field experiment; small ensembles
  - Varied physics, initial / boundary data
  - Domain: ~200x200x50
  - Cores: 256, for 1-3 hr
  - I/O: modest (5 min)
  - Cloud applicability: *high*

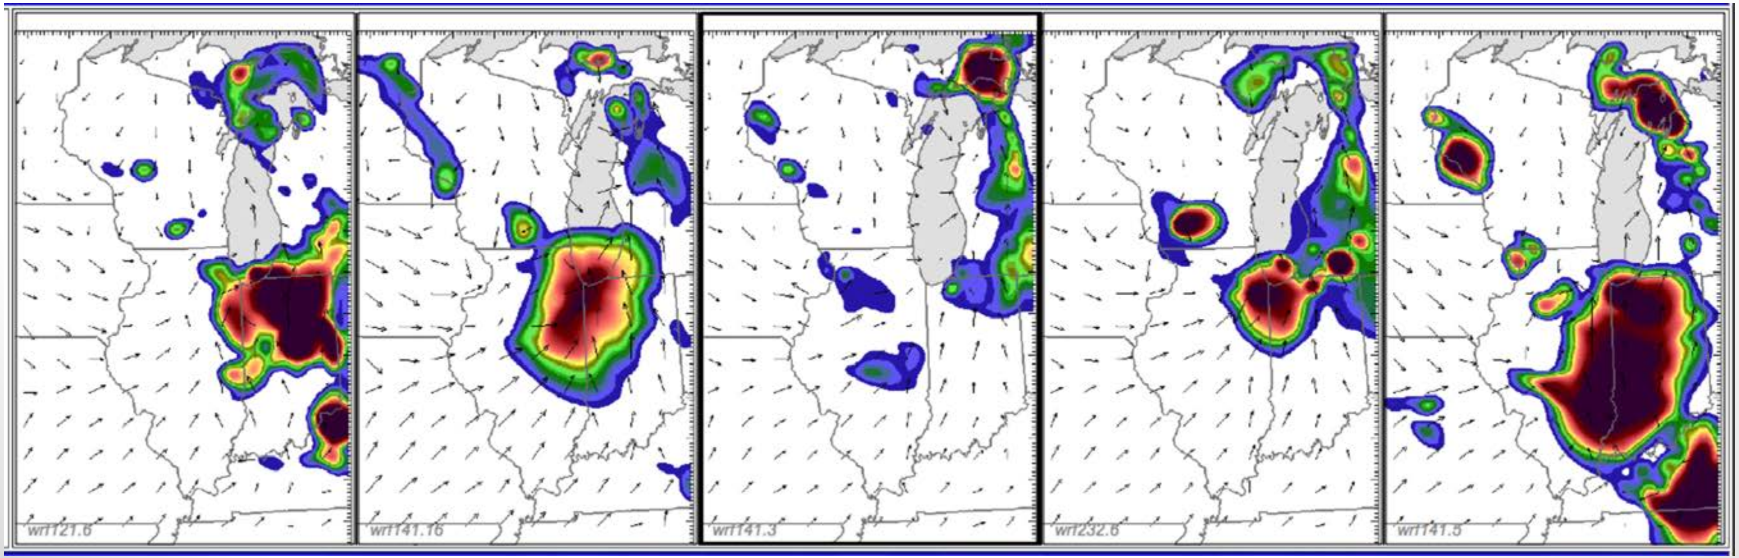

Classroom HPC use

- *Numerical fluid dynamics*
  - Graduate class, Univ. Illinois
  - Taught numerical methods and fluid modeling
  - Exposure to HPC file systems, batch system, linux OS
  - Ramping up of computing need as problems > 3D nonlinear code
  - Used Stampede @ TACC
  - Cloud applicability: *high*

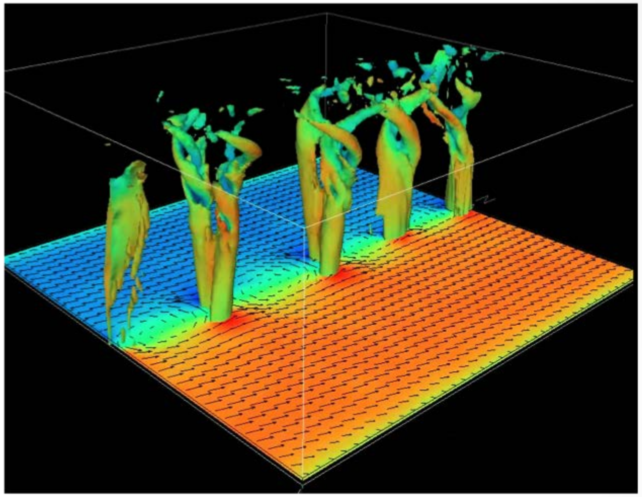

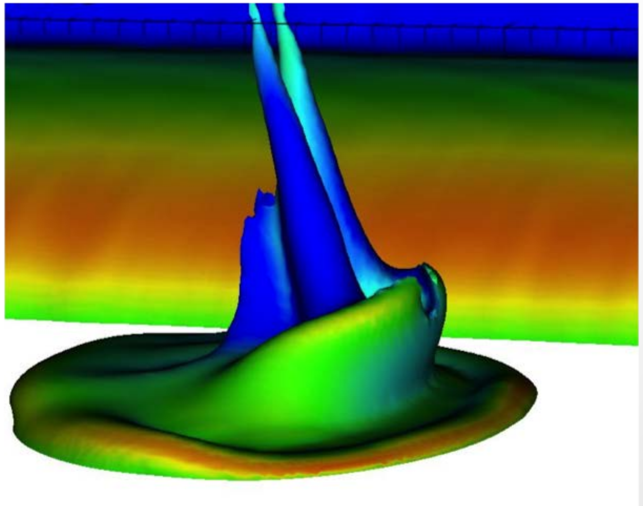

Clear*est* use cases?

- For cloud computing -
  - class use for quick turnaround, small problem sizes
  - canned app configurations applied with minimal changes
  - closely-coupled workflows with well developed tools
    - esp. when computing *is* close to the data!
- Traditional roles for traditional "big iron"
  - large(r) core counts: 512 and more
  - problems requiring very large data storage
  - parameter sweeps which may be revisited later

Closing thoughts: *a mix*

- **Big iron**
  - Very high performance, mpi/hybrid, fast networks, parallel FS
  - Reproduceability *should* be easy
    - Common hardware throughout each HPC
  - Large rotating disk + mass store accompany the big iron
    - But: HPCs have limited lifetime (transition @ TACC: good!)
- **Cloud computing**
  - doesn't "go away"
  - terrific opportunity to broaden acceptance, use, and tools
  - potentially speed up time-to-solution esp. for newcomers
    - less time learning/doing, more time understanding
- **Research use of the cloud**
  - How will SUs be awarded? Are SUs cheaper for NSF?
  - XSEDE flexibility at present (startups, extensions, etc)
    - Will we be able to - or even need to - shift allocations?