Unlocking the Mysteries of the Most Violent Tornadoes

Leigh Orf

Cooperative Institute for Meteorological Satellite Studies
University of Wisconsin - Madison

Collaborators:

Bruce Lee, High Impact Weather Research and Consulting, LLC
Catherine Finley, St. Louis University
Robert Wilhelmson, NCSA/UIUC

Blue Waters Symposium
Sunriver, OR
May 16, 2017

PRAC: Understanding the development and evolution of violent tornadoes in supercell thunderstorms

Outline

Research goals

Why it matters

Why Blue Waters

Results: Simulations in the 24 May 2011 (El Reno, OK) environment

ZFP floating point lossy compression

Parting thoughts

Research goals

- Simulate thunderstorm phenomena at the highest spatial resolution possible in order to properly capture important physical processes, especially near the ground
- Visualize, with high fidelity, simulation data at very high temporal resolution in order to identify features for quantitative study
- Overarching goal: understand the internal workings of supercell thunderstorms well enough to significantly improve our forecasting of their behavior

Why study storms that produce high-end, long-path tornadoes?

- Large majority ($>$ 90%) of fatalities occur in EF4-EF5 strength tornadoes
- Long-path tornadoes offer potential for most damage
- Despite advances in observational and computational meteorology, why some supercells produce no / weak / strong tornado is not currently understood

What it takes to make this all work

1. A good model - CM1
    - Developed at NCAR, contains about 10,000 lines of Fortran
    - Can utilize hundreds of thousands of processing elements (CPUs) efficiently
2. A good set of initial atmospheric conditions
    - We chose the environment adjacent to the May 24, 2011 El Reno Oklahoma supercell
3. Good hardware (!)
4. Good software
    - Lots of custom code for I/O management (eventually to be shared with the community)
    - VisIt, Vapor for visualization
5. A bit of luck!

Why Blue Waters

1. The scale of the problem demands “bleeding edge” HPC
2. 24 May 2011 with 30 meter isotropic inner mesh:
    - $2200 \times 2200 \times 380$ = **1.8 billion grid points**
    - $120 \times 120 \times 20$ km domain
    - 30 m isotropic inner mesh ($60 \times 60 \times 10$ km)
3. 24 May 2011 with 20 meter isotropic inner mesh:
    - $3600 \times 3600 \times 600$ = **7.8 billion grid points**
    - $120 \times 120 \times 20$ km domain
    - 20 m isotropic inner mesh ($60 \times 60 \times 10$ km)
4. We’ve created a total of $\sim$3 PB of data over the years, with around 500 TB of EF5 tornadic storm data

Key challenges

1. Getting the model to create the desired storm
    - The type of tornado we desire to study is the least common
    - It took many attempts before we managed to get a supercell that produces a long track EF5
2. Handling the data load
    - In order to meet our research objectives, data is saved frequently (every model second) for visualization and analysis
    - I've spent a lot more time over the past decade wrestling with the data problem than doing tornado science
3. Visualizing data in a way that elucidates important structure and allows for comparisons to field observations

Blue Waters team contributions

- Rob Sisneros (NCSA) helped us out with some crucial bits (VisIt, ZFP)
- Gerd Heber (HDF, ZFP) via PAID program
- help+bw always useful
- THANK YOU to Blue Waters staff!!

Accomplishments/broader impacts

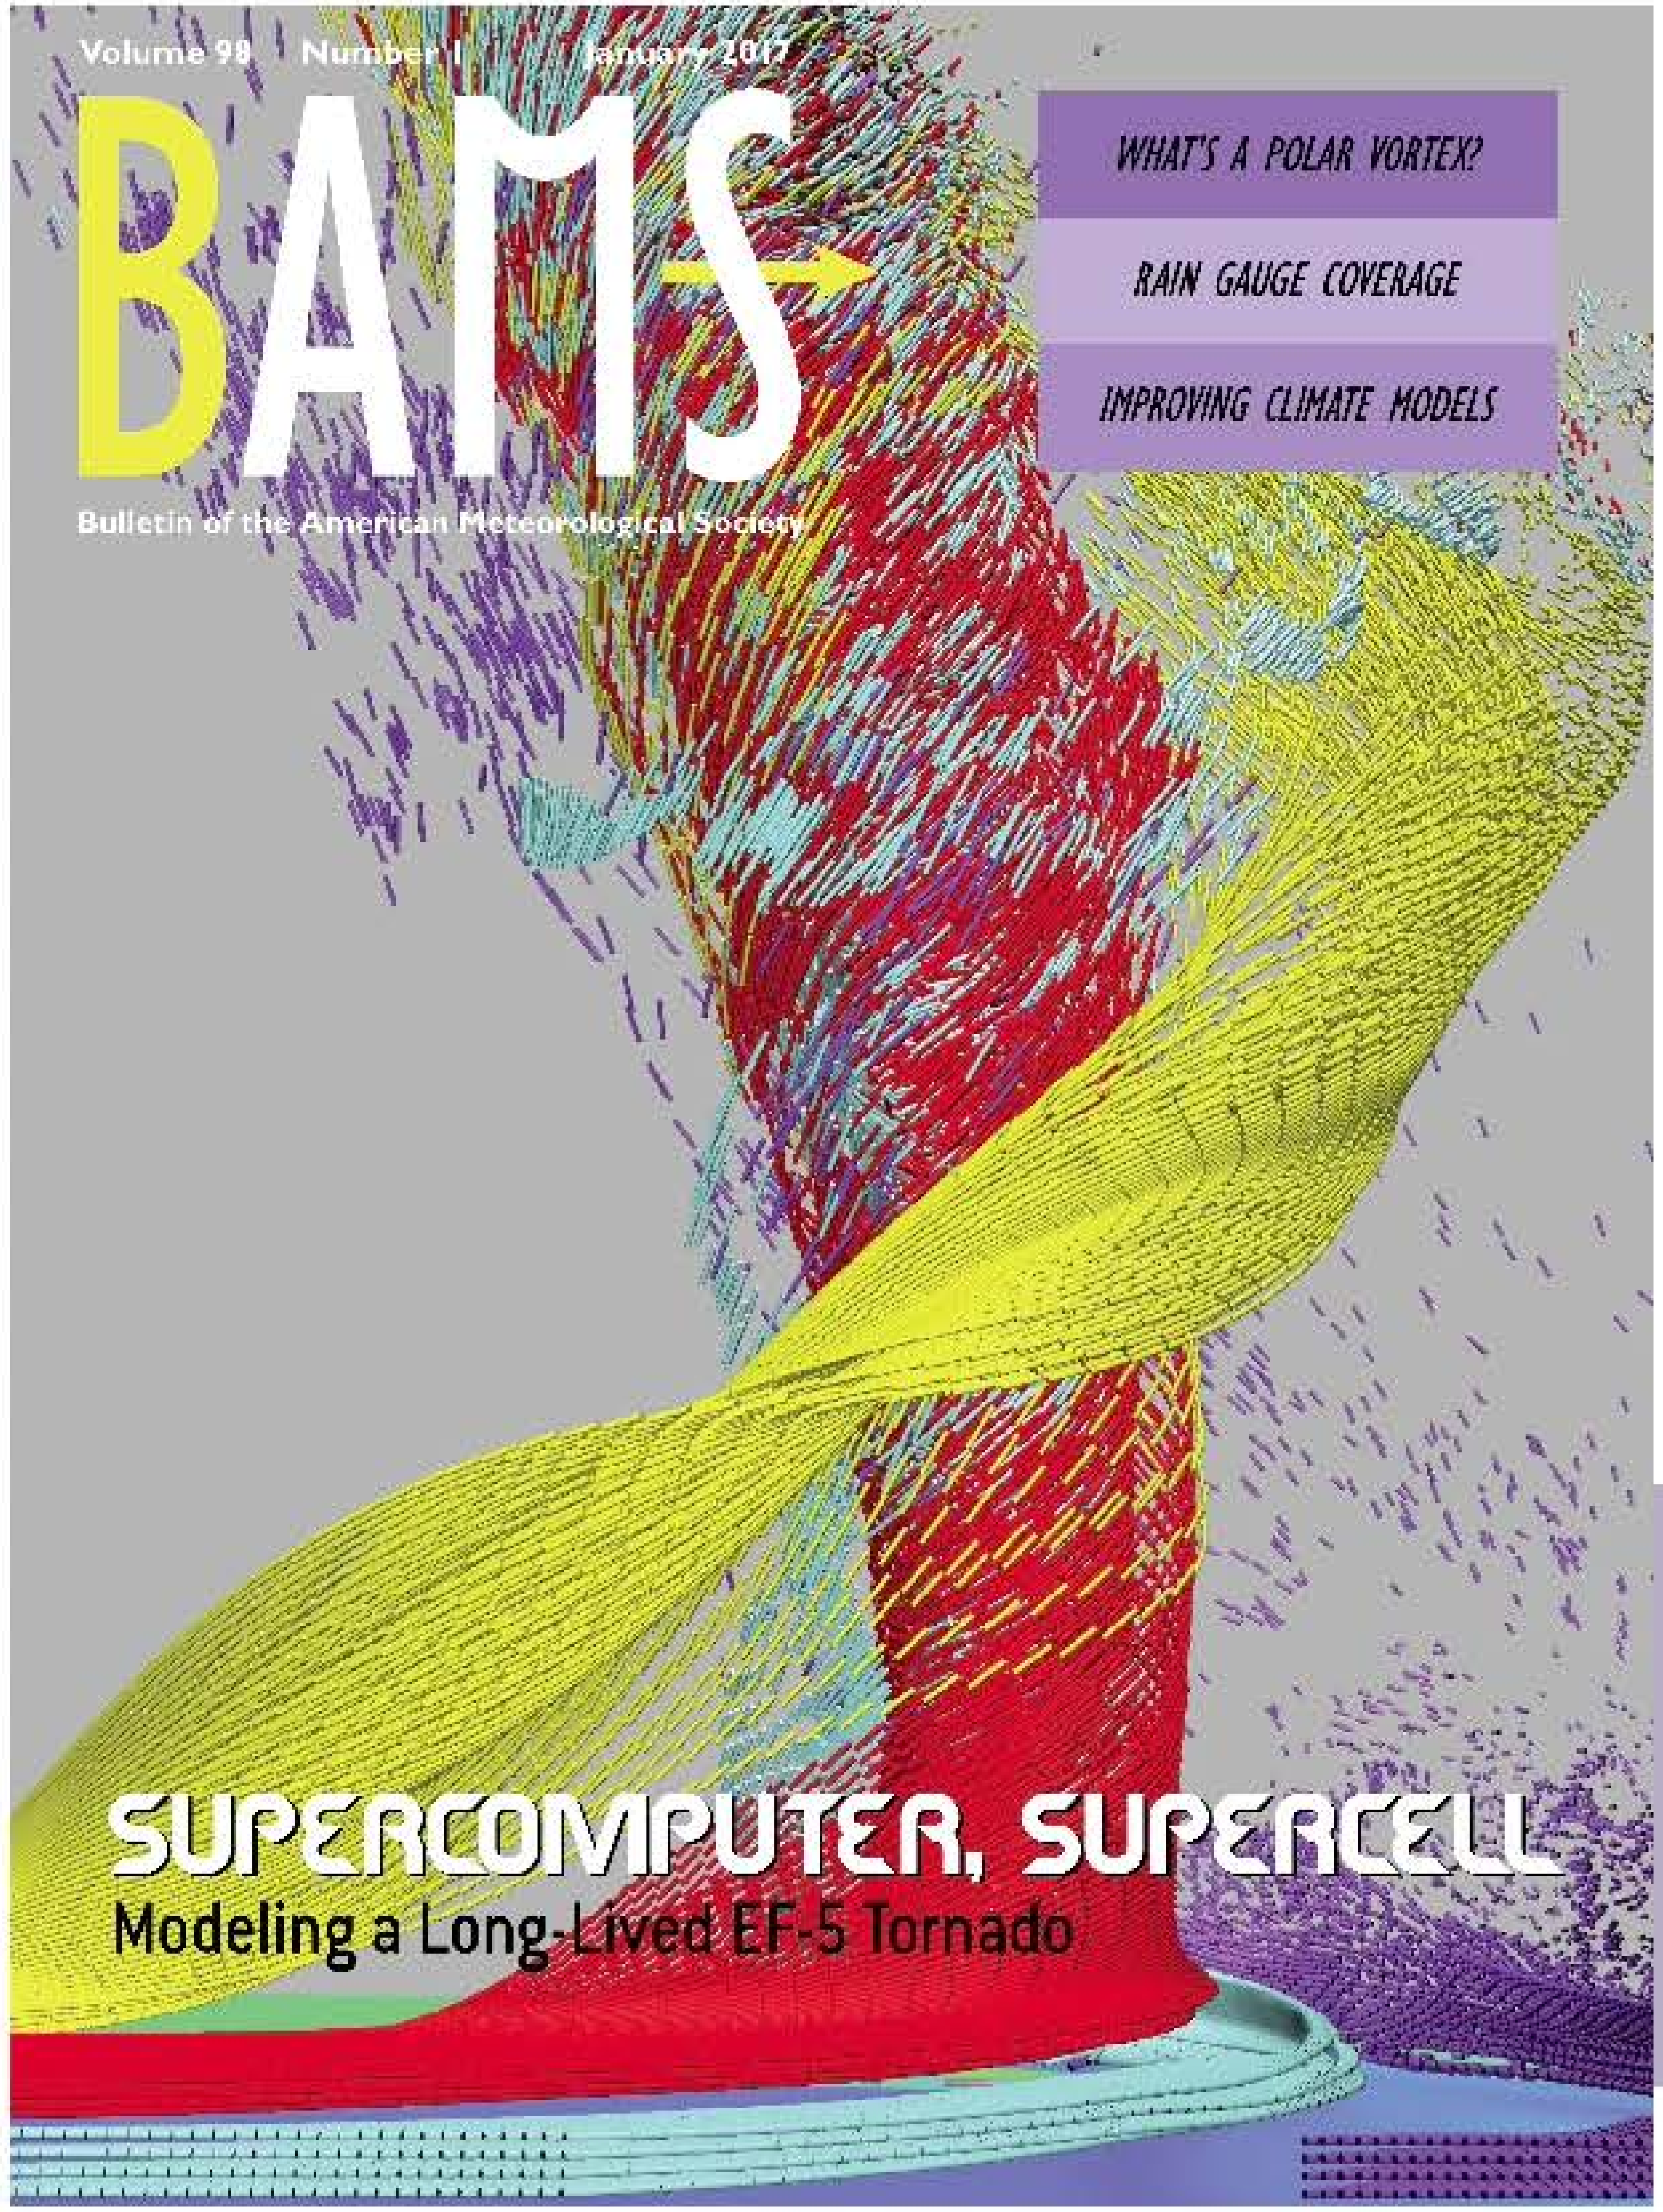

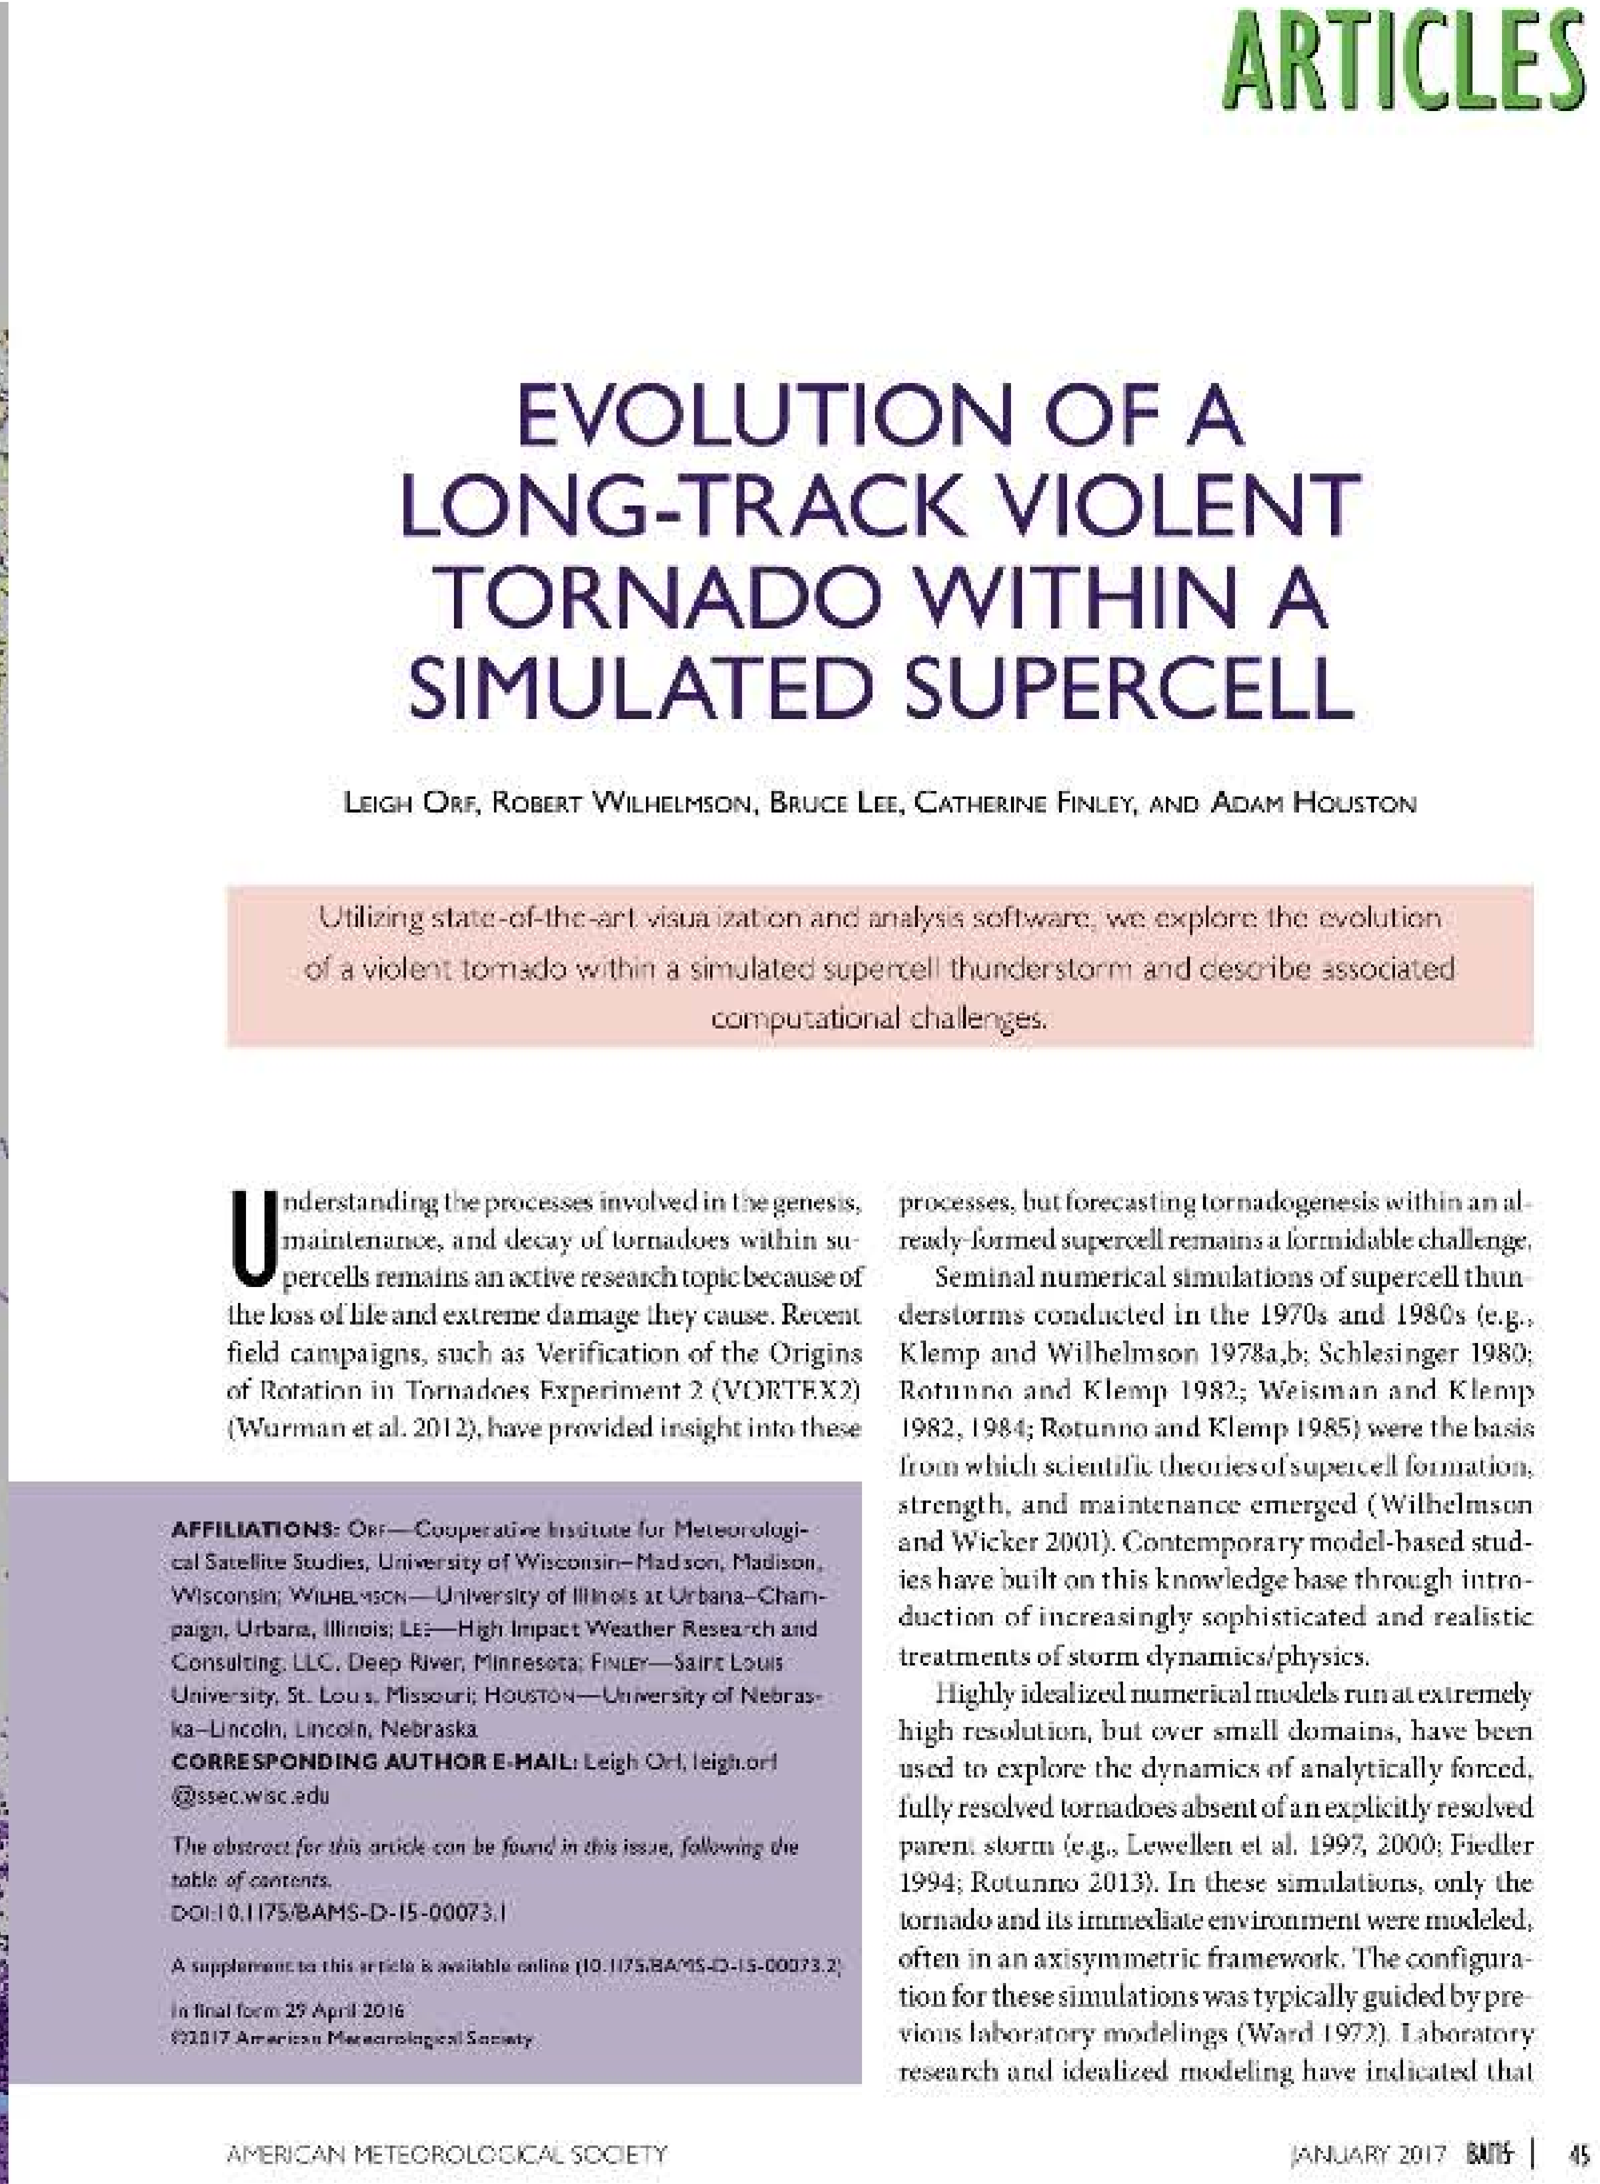

ARTICLES

EVOLUTION OF A LONG-TRACK VIOLENT TORNADO WITHIN A SIMULATED SUPERCELL

Leigh Orf, Robert Wilhelmson, Bruce Lee, Catherine Finley, and Adam Houston

Utilizing state-of-the-art visualization and analysis software, we explore the evolution of a violent tornado within a simulated supercell thunderstorm and describe associated computational challenges.

Understanding the processes involved in the genesis, maintenance, and decay of tornadoes within supercells remains an active research topic because of the loss of life and extreme damage they cause. Recent field campaigns, such as Verification of the Origins of Rotation in Tornadoes Experiment 2 (VORTEX2) (Wurman et al. 2012), have provided insight into these processes, but forecasting tornadogenesis within an already-formed supercell remains a formidable challenge.

Seminal numerical simulations of supercell thunderstorms conducted in the 1970s and 1980s (e.g., Klemp and Wilhelmson 1978a,b; Schlesinger 1980; Rotunno and Klemp 1982; Weisman and Klemp 1982, 1984; Rotunno and Klemp 1985) were the basis from which scientific theories of supercell formation, strength, and maintenance emerged (Wilhelmson and Wicker 2001). Contemporary model-based studies have built on this knowledge base through introduction of increasingly sophisticated and realistic treatments of storm dynamics/physics.

Highly idealized numerical models run at extremely high resolution, but over small domains, have been used to explore the dynamics of analytically forced, fully resolved tornadoes absent of an explicitly resolved parent storm (e.g., Lewellen et al. 1997, 2000; Fiedler 1994; Rotunno 2013). In these simulations, only the tornado and its immediate environment were modeled, often in an axisymmetric framework. The configuration for these simulations was typically guided by previous laboratory modelings (Ward 1972). Laboratory research and idealized modeling have indicated that

AFFILIATIONS: Orf—Cooperative Institute for Meteorological Satellite Studies, University of Wisconsin–Madison, Madison, Wisconsin; Wilhelmson—University of Illinois at Urbana–Champaign, Urbana, Illinois; Lee—High Impact Weather Research and Consulting, LLC, Deep River, Minnesota; Finley—Saint Louis University, St. Louis, Missouri; Houston—University of Nebraska–Lincoln, Lincoln, Nebraska
CORRESPONDING AUTHOR E-MAIL: Leigh Orf, leigh.orf@ssec.wisc.edu

The abstract for this article can be found in this issue, following the table of contents.
DOI:10.1175/BAMS-D-15-00073.1

A supplement to this article is available online (10.1175/BAMS-D-15-00073.2)

In final form 29 April 2016

AMERICAN METEOROLOGICAL SOCIETY

JANUARY 2017 BAMS | 45

- CADENS documentary
- Three journal articles, 1 MS thesis (SLU)
- Lots of local/state/national/international PR following UW press release (Weather Channel, PBS Newshour, etc.)
- Many invited presentations
- Two graduate students coming to UW to join research team

Overview of 24 May 2011 30 meter simulation

- EF5-strength tornado exists from 1-2 hours (depending on specific simulation)
- Non-tornadic vortices are found in abundance along the boundary between the cooler, denser, storm-generated near-surface air and the environmental air
- Vortices “accumulate” (and merge) near the inflection point beneath the updraft, between the two major sources of downdraft air (RFD, FFD)
- In addition to these small-scale vortex mergers, we have identified a feature we call the streamwise vorticity current (SVC) that transports baroclinically-generated streamwise horizontal vorticity into the vertical

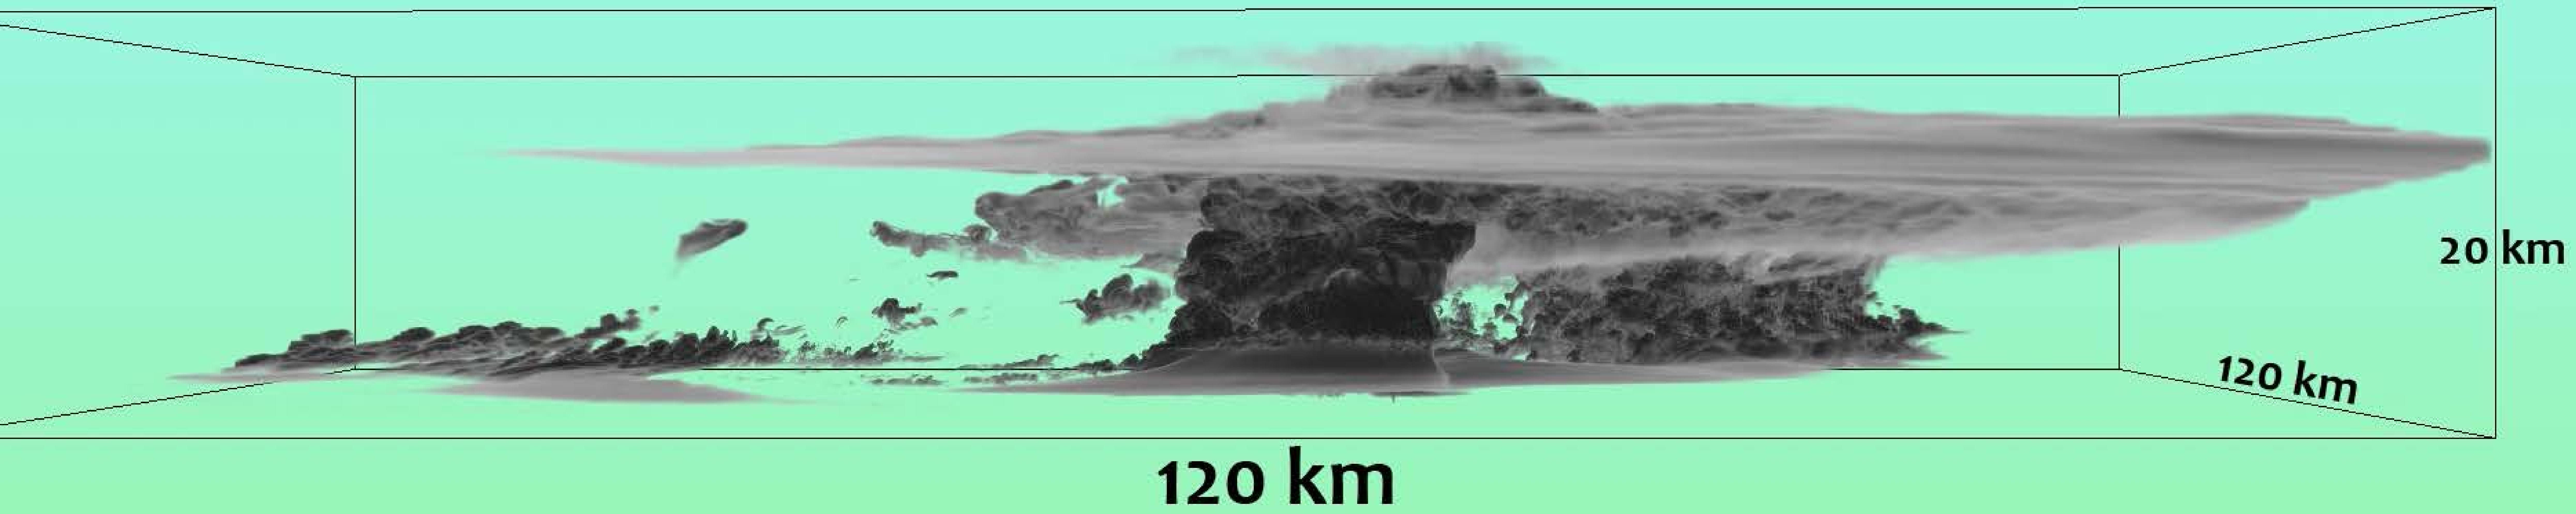
20 km
120 km
120 km


N
sfc
θ'_ρ
-8 -7 -6 -5 -4 -3 -2 -1 0 1 2
1km
dBZ
-30 -20 -10 0 10 20 30 40 50
1km
w
-10 0 10 20 30 40 50
looking west
N
2:02:18


DB: fs8b-b.cm1visit
Time:3976
Tornado
Z
Y
X
-2
-4
0
-2
-4
-6
looking SE
1:06:16


-2
6
-8
-2
-4

Peter Lindstrom's ZFP compression and HDF

- One of the biggest challenges of this work is the data load
- We have recently begun to utilize ZFP, which is available as a module on Blue Waters
- ZFP has allowed us to reduce our data load by around **30x** compared to lossless gzip
- Question to ask yourself: *Do I need 64 (or 32) bits of precision for all my output? What if I could save data with only the amount of accuracy I need?*
- ZFP is trivially easy to use
- Huge bonus: You can specify the RMS accuracy of your data, e.g., 0.1 Kelvins, 1 dBZ, etc.

ZFP compression and HDF

Super easy!

```
module module_orfio
use hdf5
use h5zzfp_props_f
...
ierr = H5Z_zfp_initialize()
...
call h5pcreate_f(H5P_DATASET_CREATE_F,chunk_id,ierr)
call h5pset_chunk_f(chunk_id,rank,chunkdims,ierr)
ierr = h5pset_zfp_accuracy(chunk_id,accuracy)
call h5dcreate_f(f_id,trim(varname),H5T_NATIVE_REAL,dspace_id, \
      dset_id,ierr,dcpl_id=chunk_id)
call h5dwrite_f(...)
...
```

ZFP compression and HDF

- ZFP requires serial HDF
- All our 3D data is saved on a per-node basis, so this is fine (and we use the core HDF driver to grow our HDF files in memory)
- Performance of ZFP compression is very good (was written to be very fast)
- Some I/O performance issues remain, but I do not believe they are ZFP related

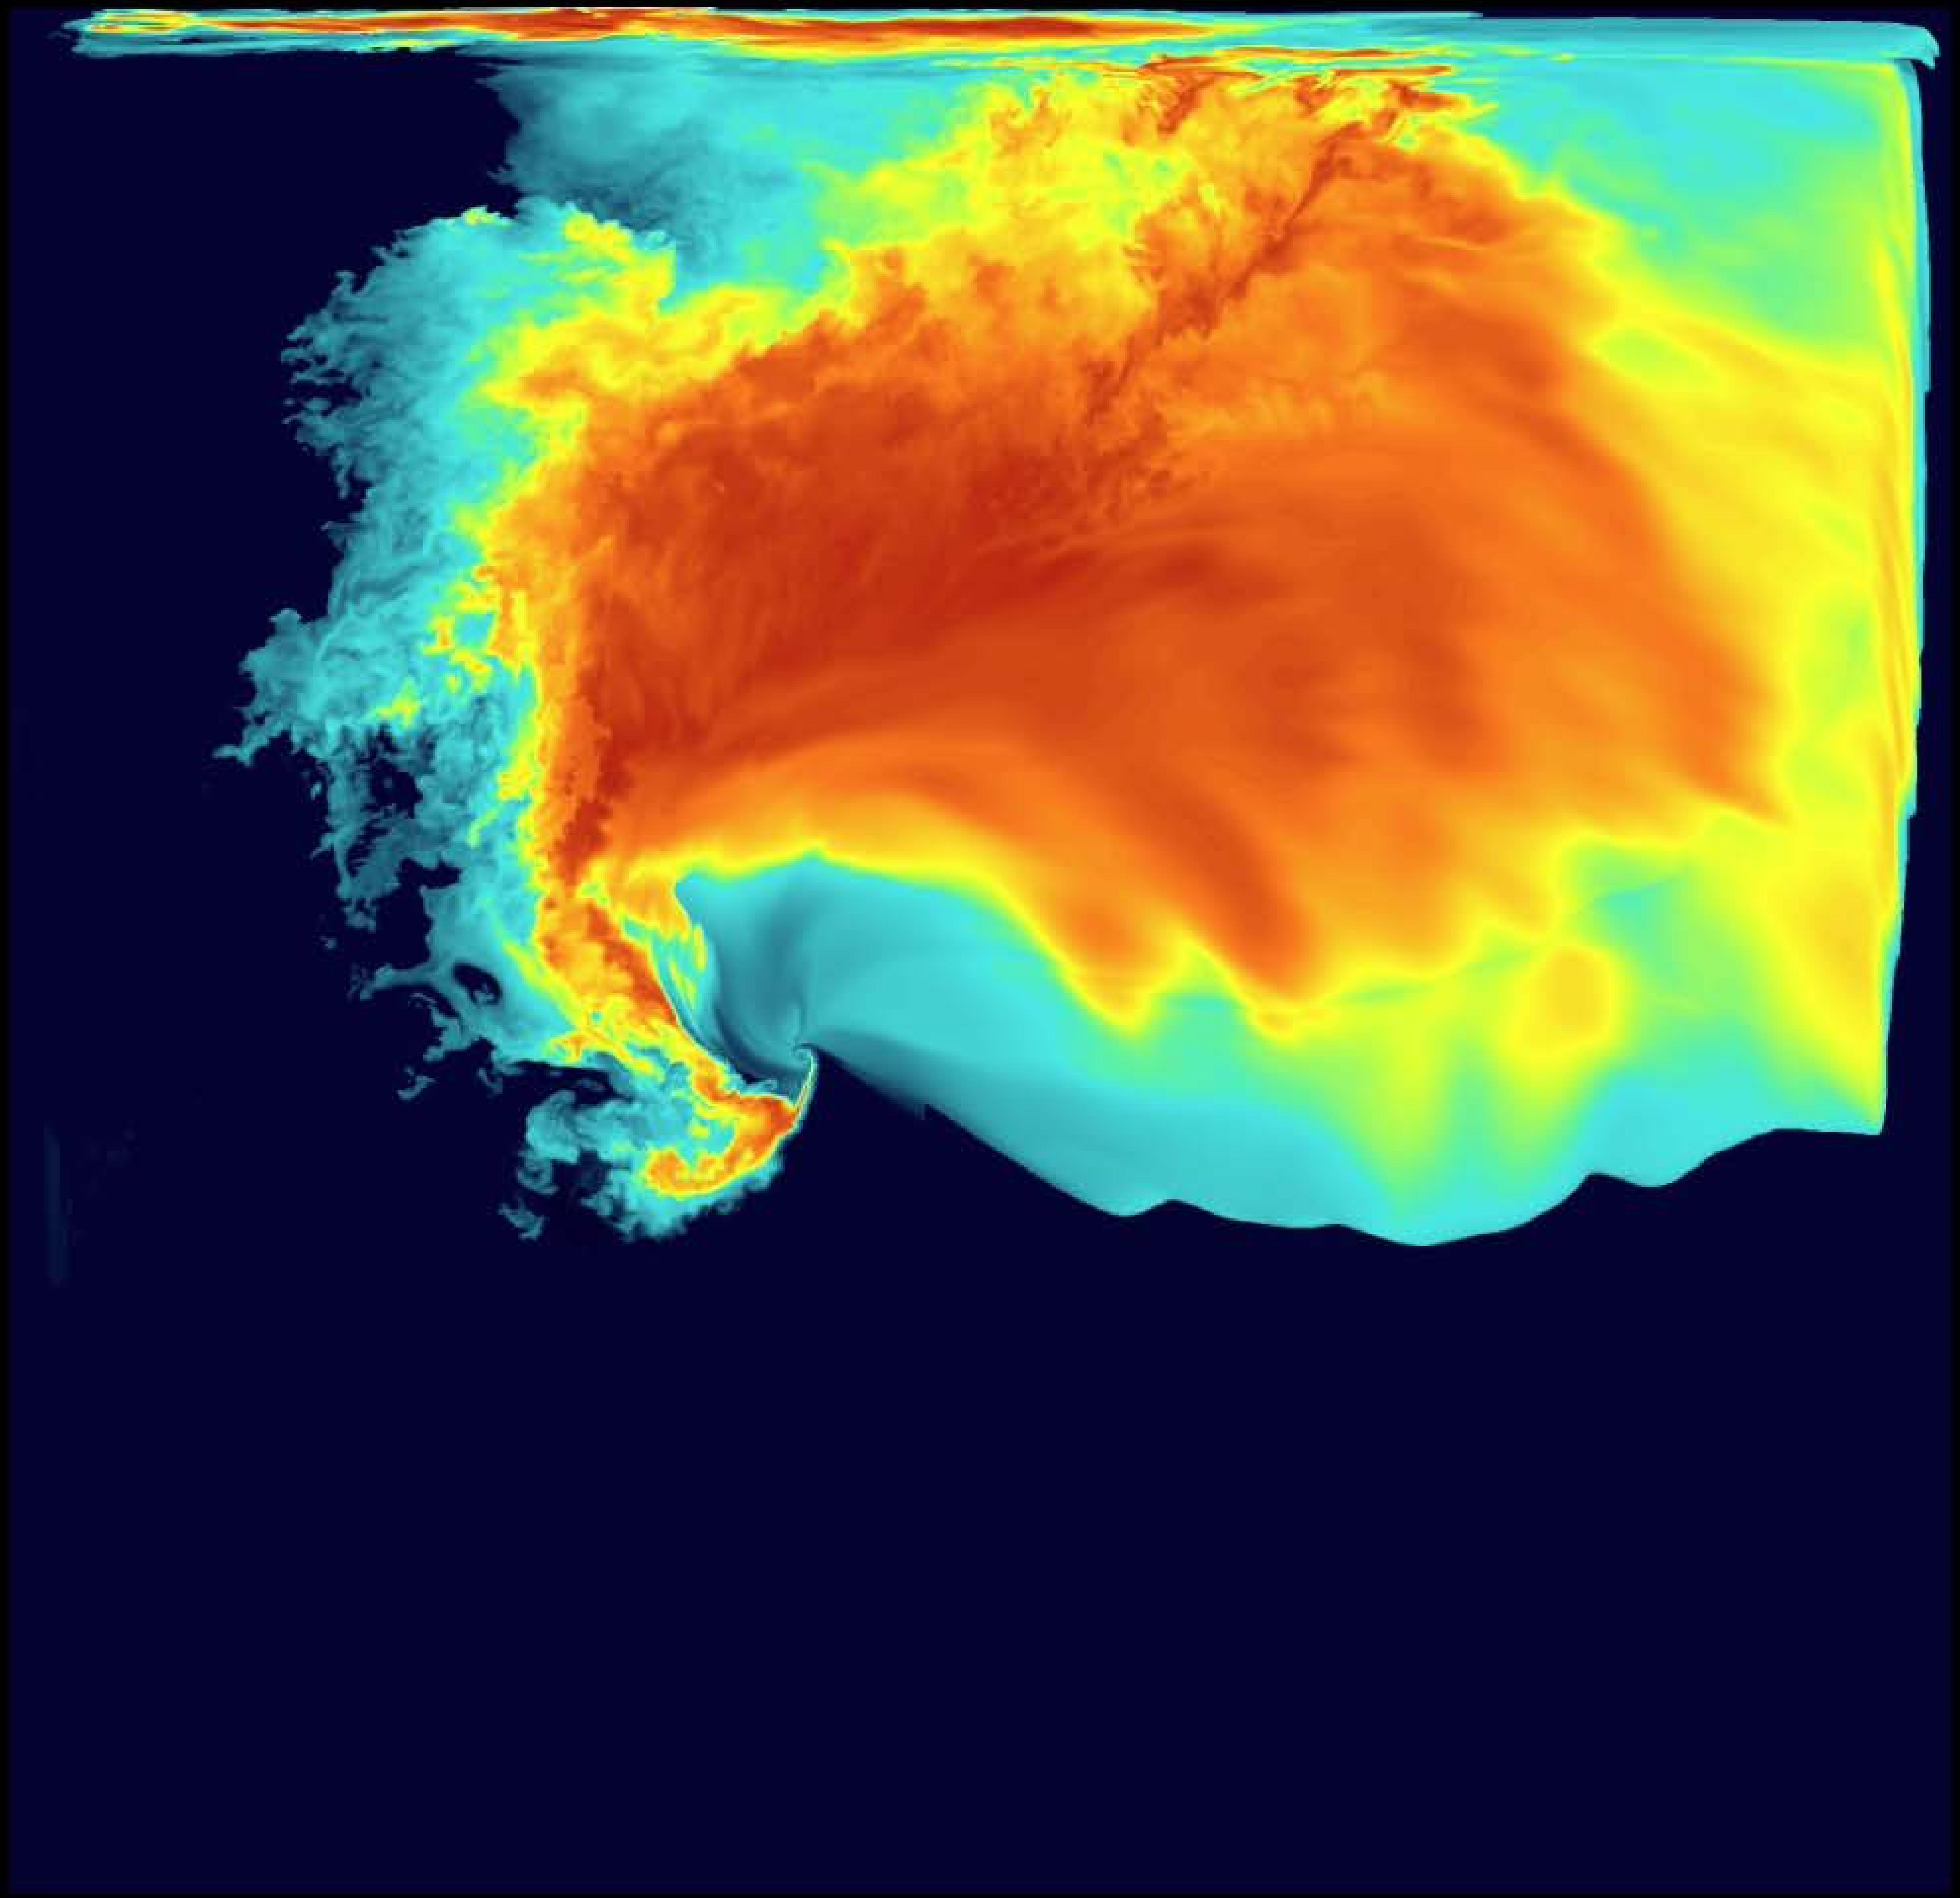

Radar reflectivity (dBZ) 1 km AGL

Next slide will show compression ratios for a constant accuracy parameter for this variable (full 3D fields at each node location)...

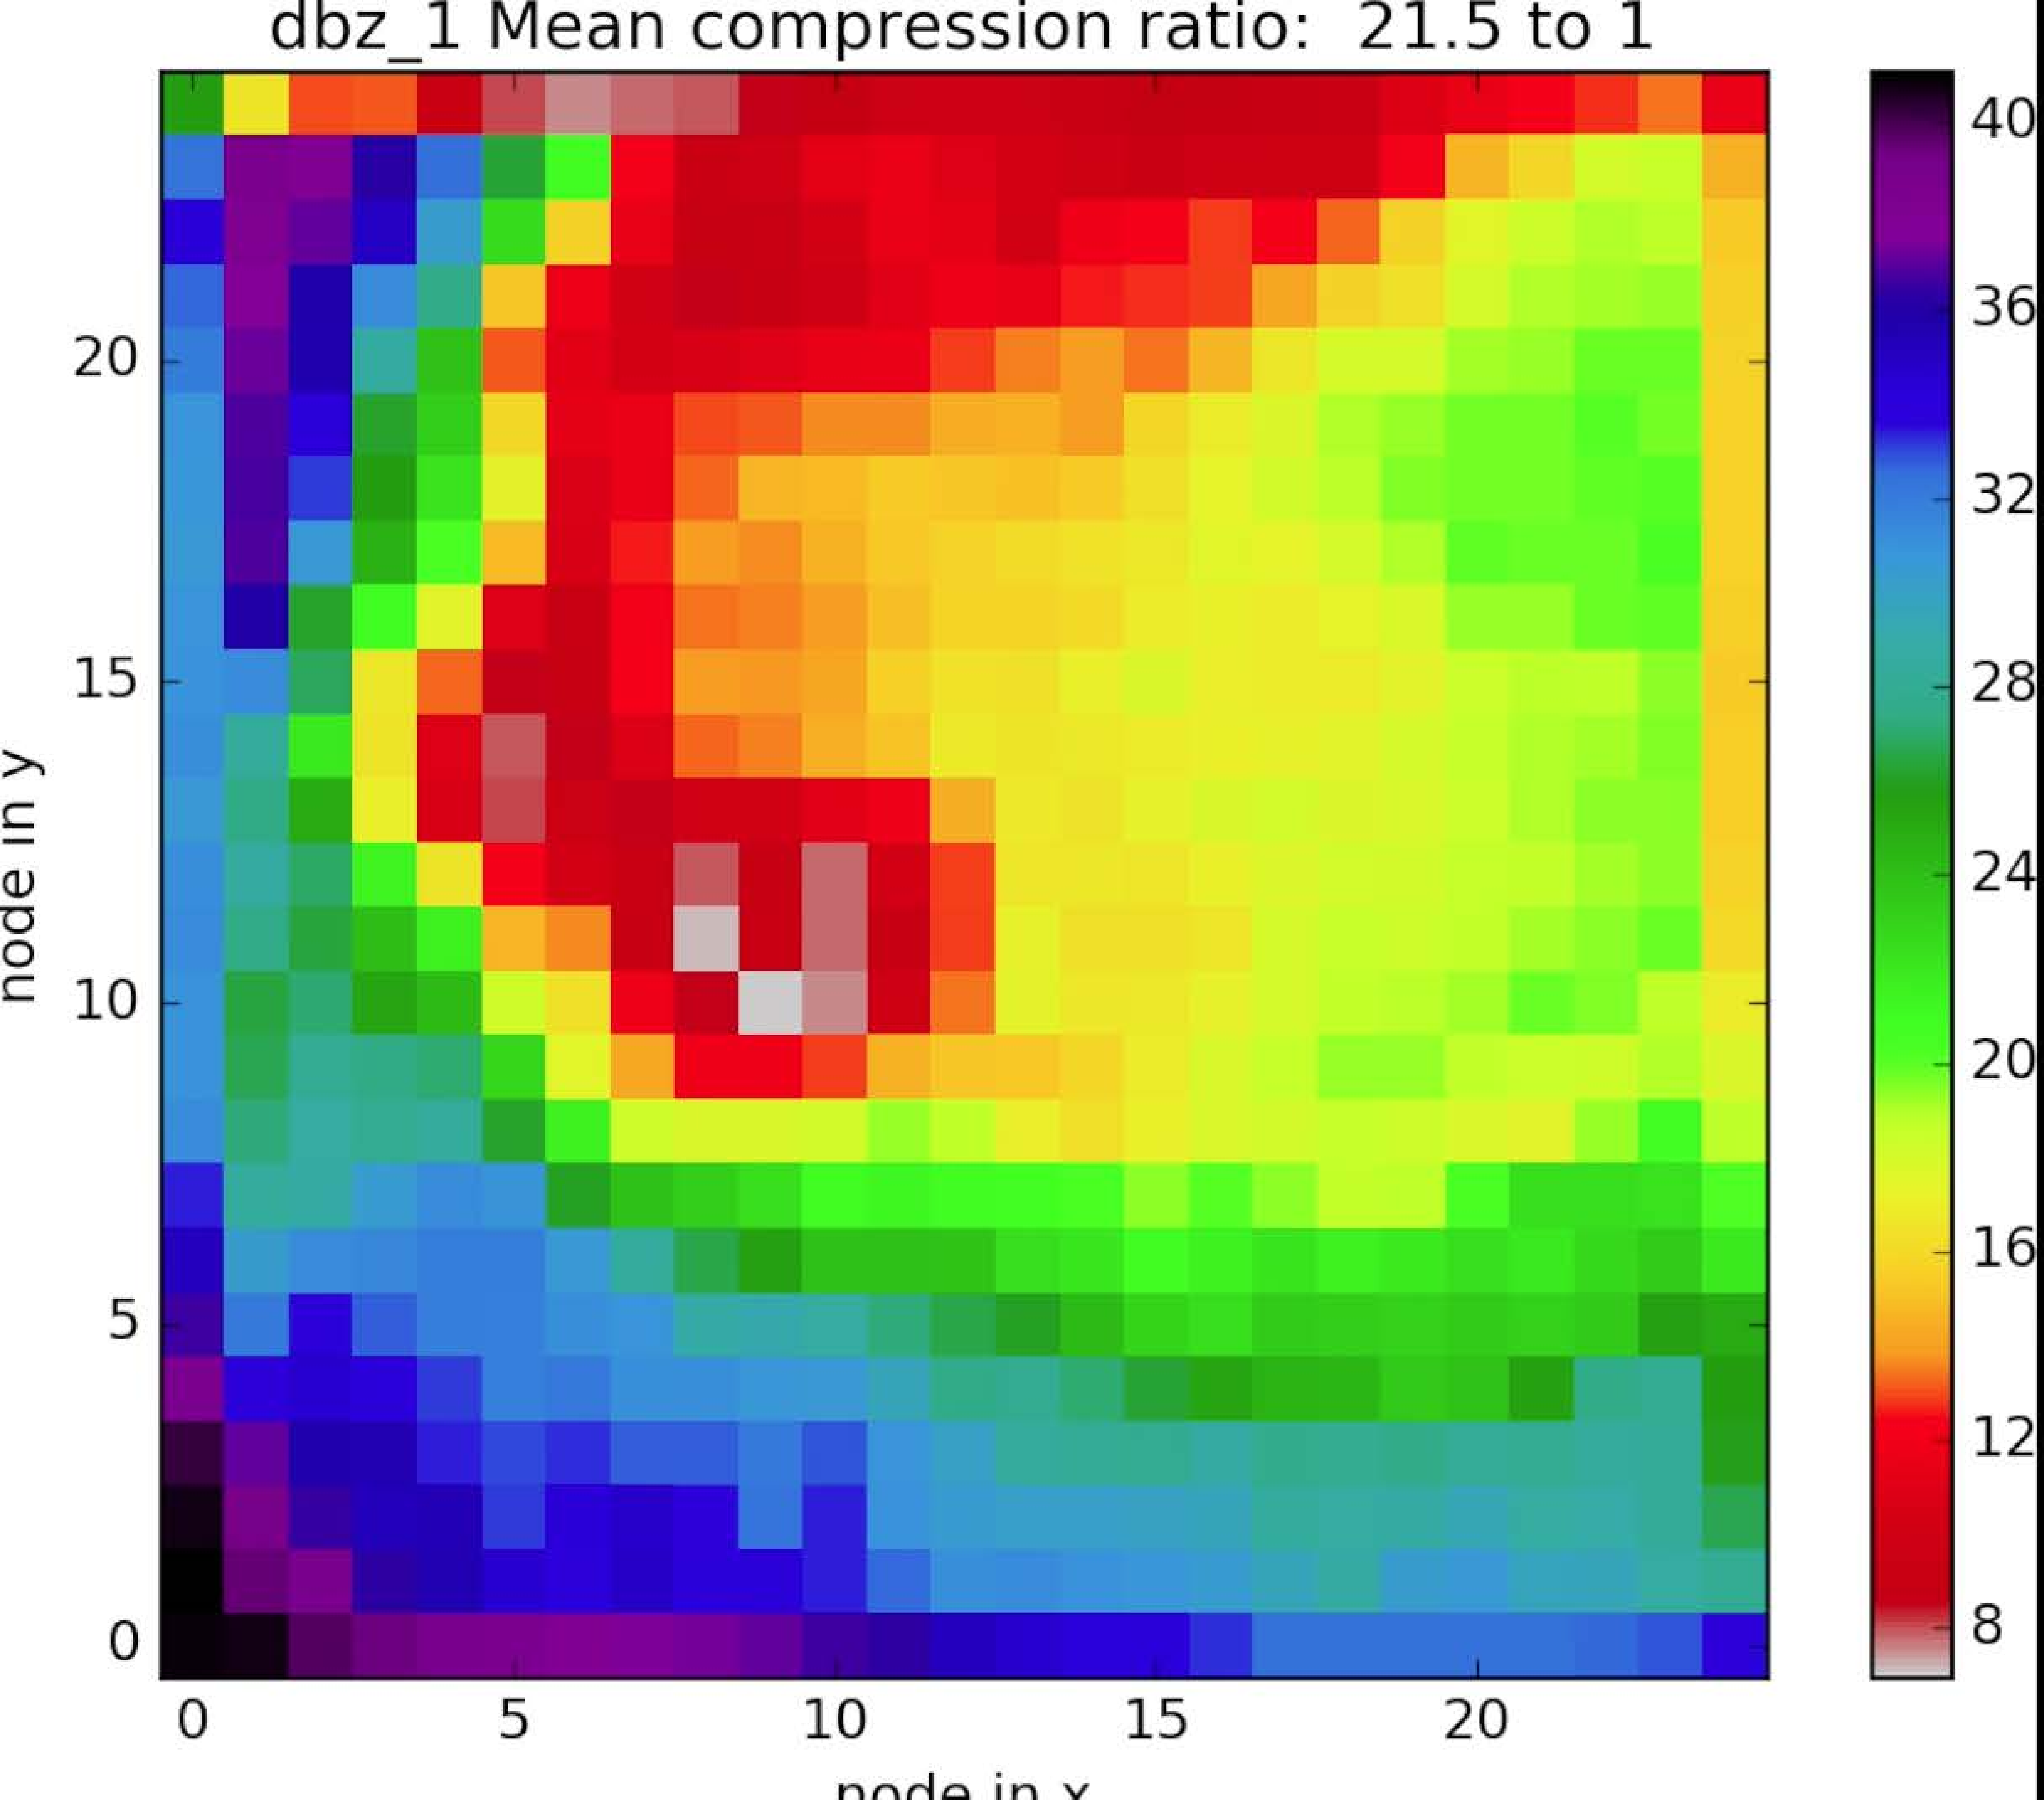

dBZ ranges from -40 to 70

Accuracy=1.0

Domain-wide compression ratio 22:1

Lots of variation throughout domain, with the "worst" compression ratios where most of the action is

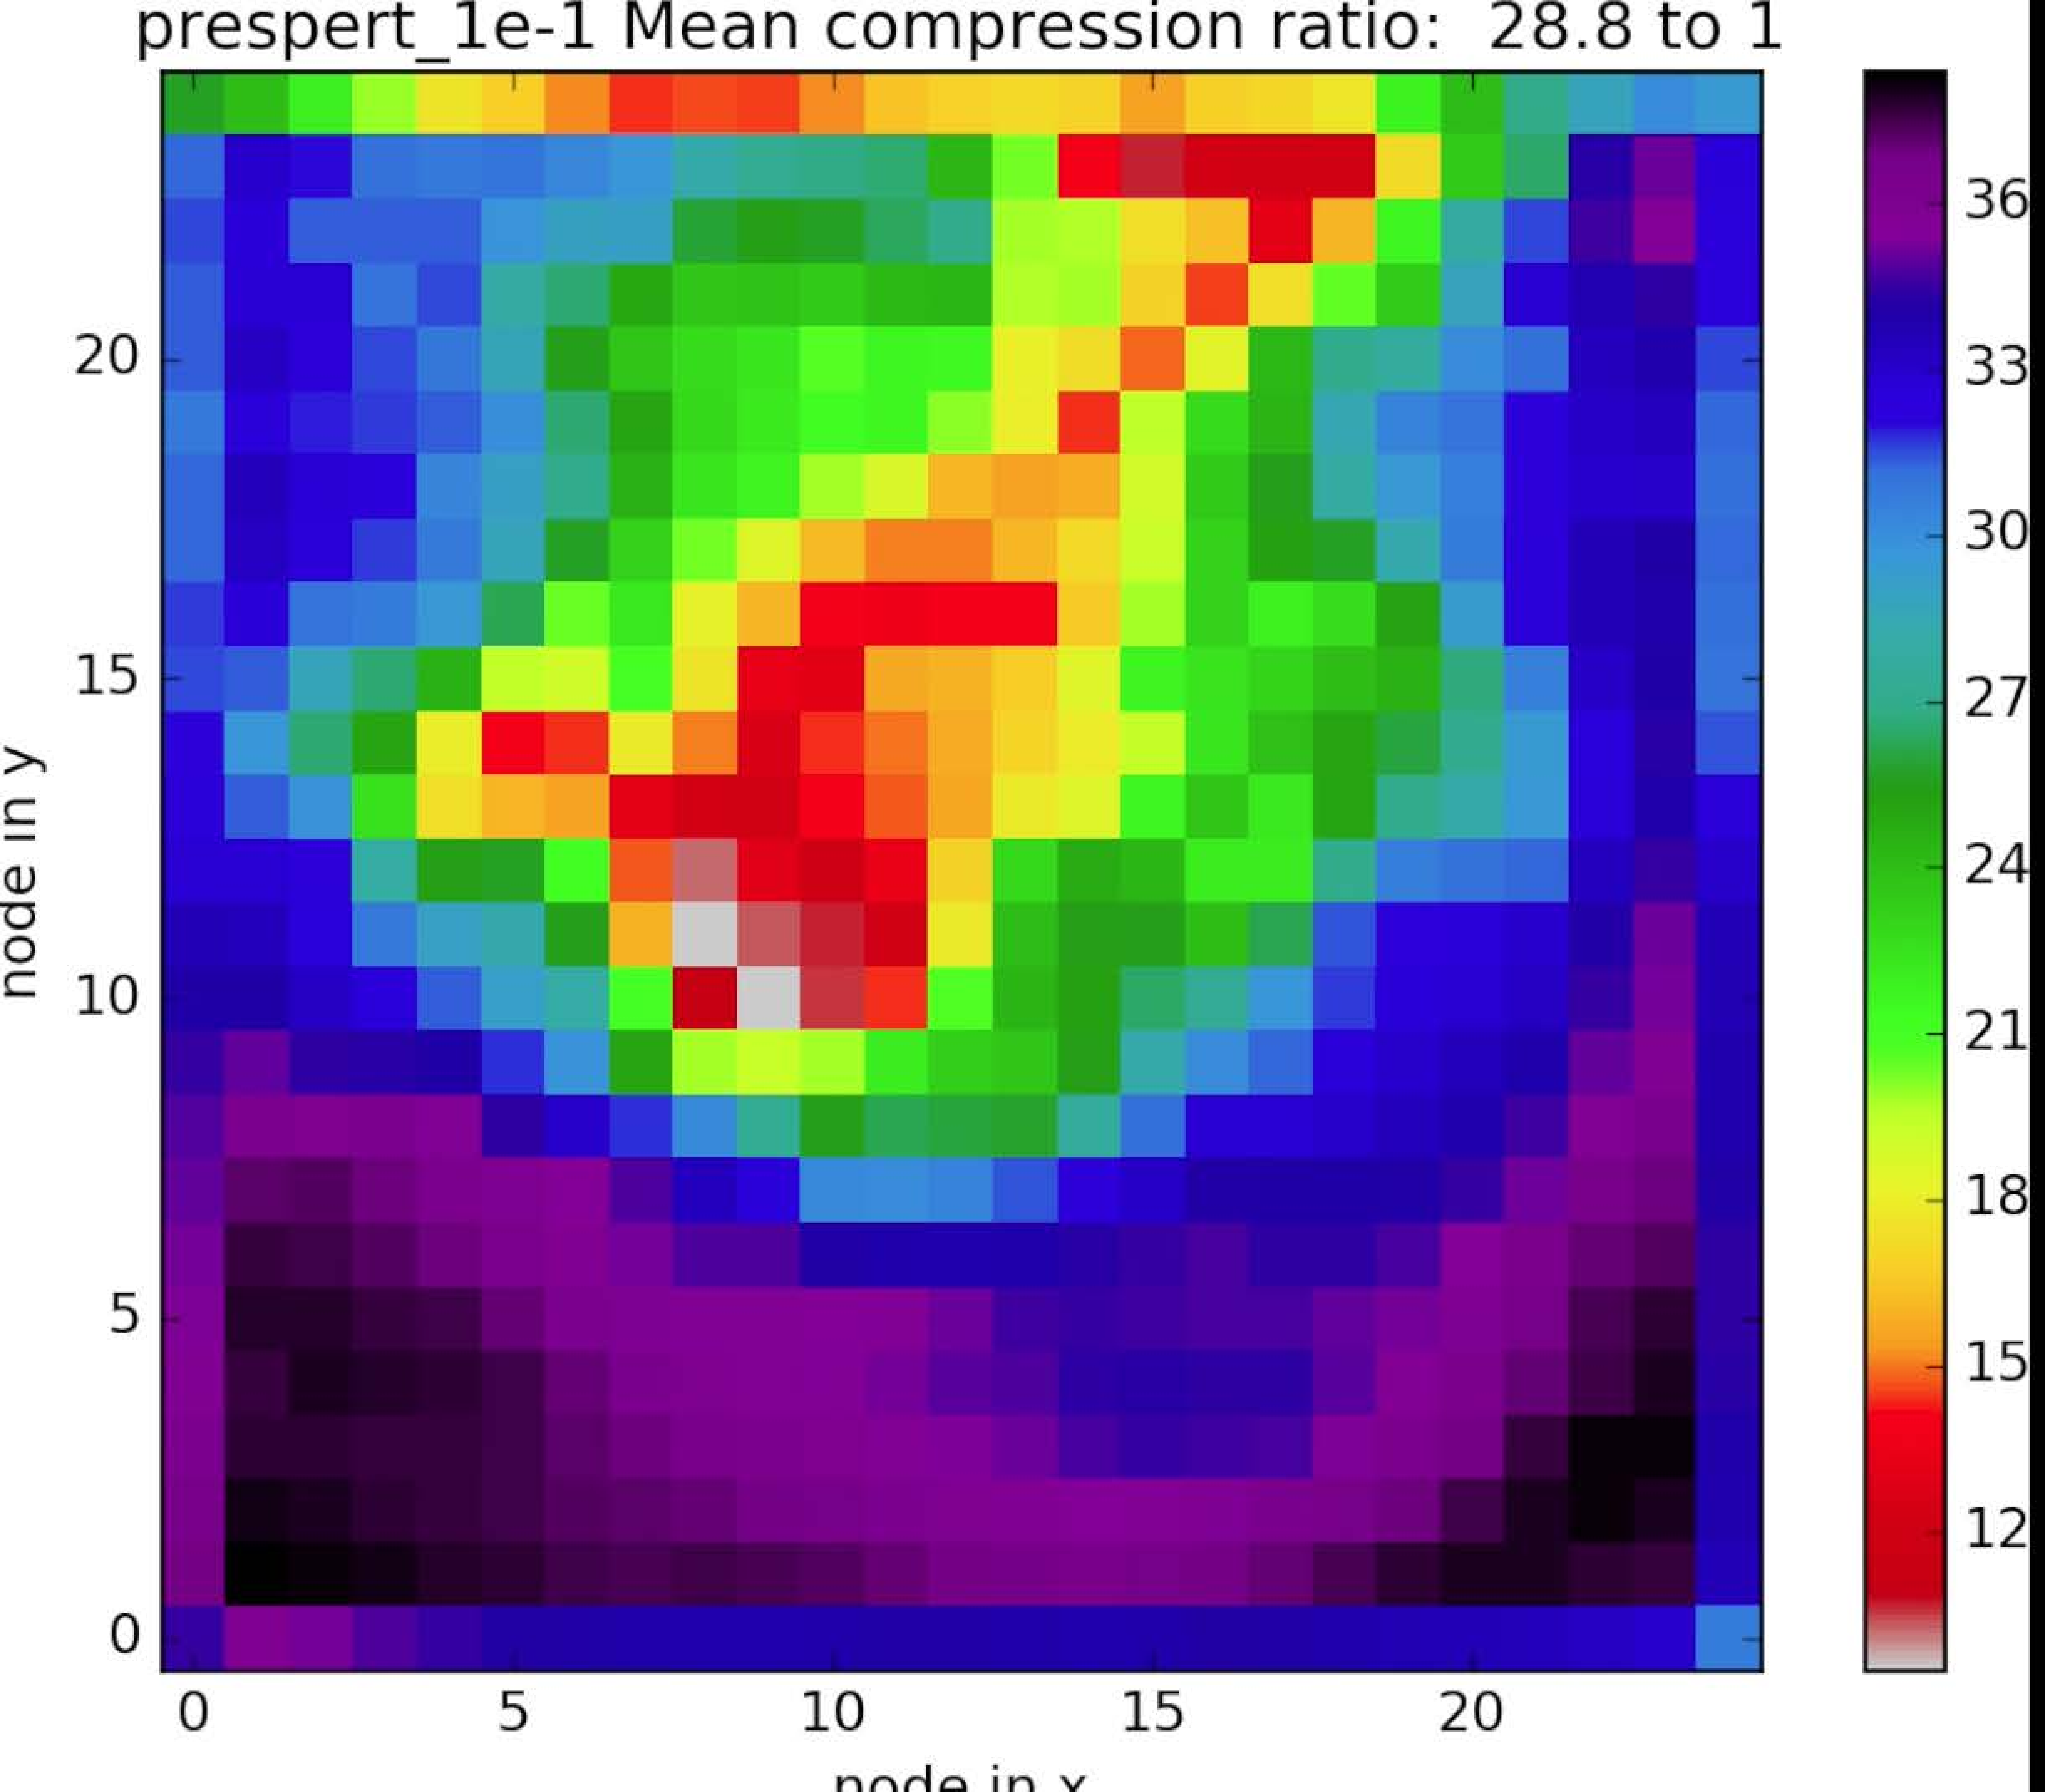

pressure perturbation ranges from -200 to 50 hPa

Accuracy=0.1 hPa

Domain-wide compression ratio 29:1

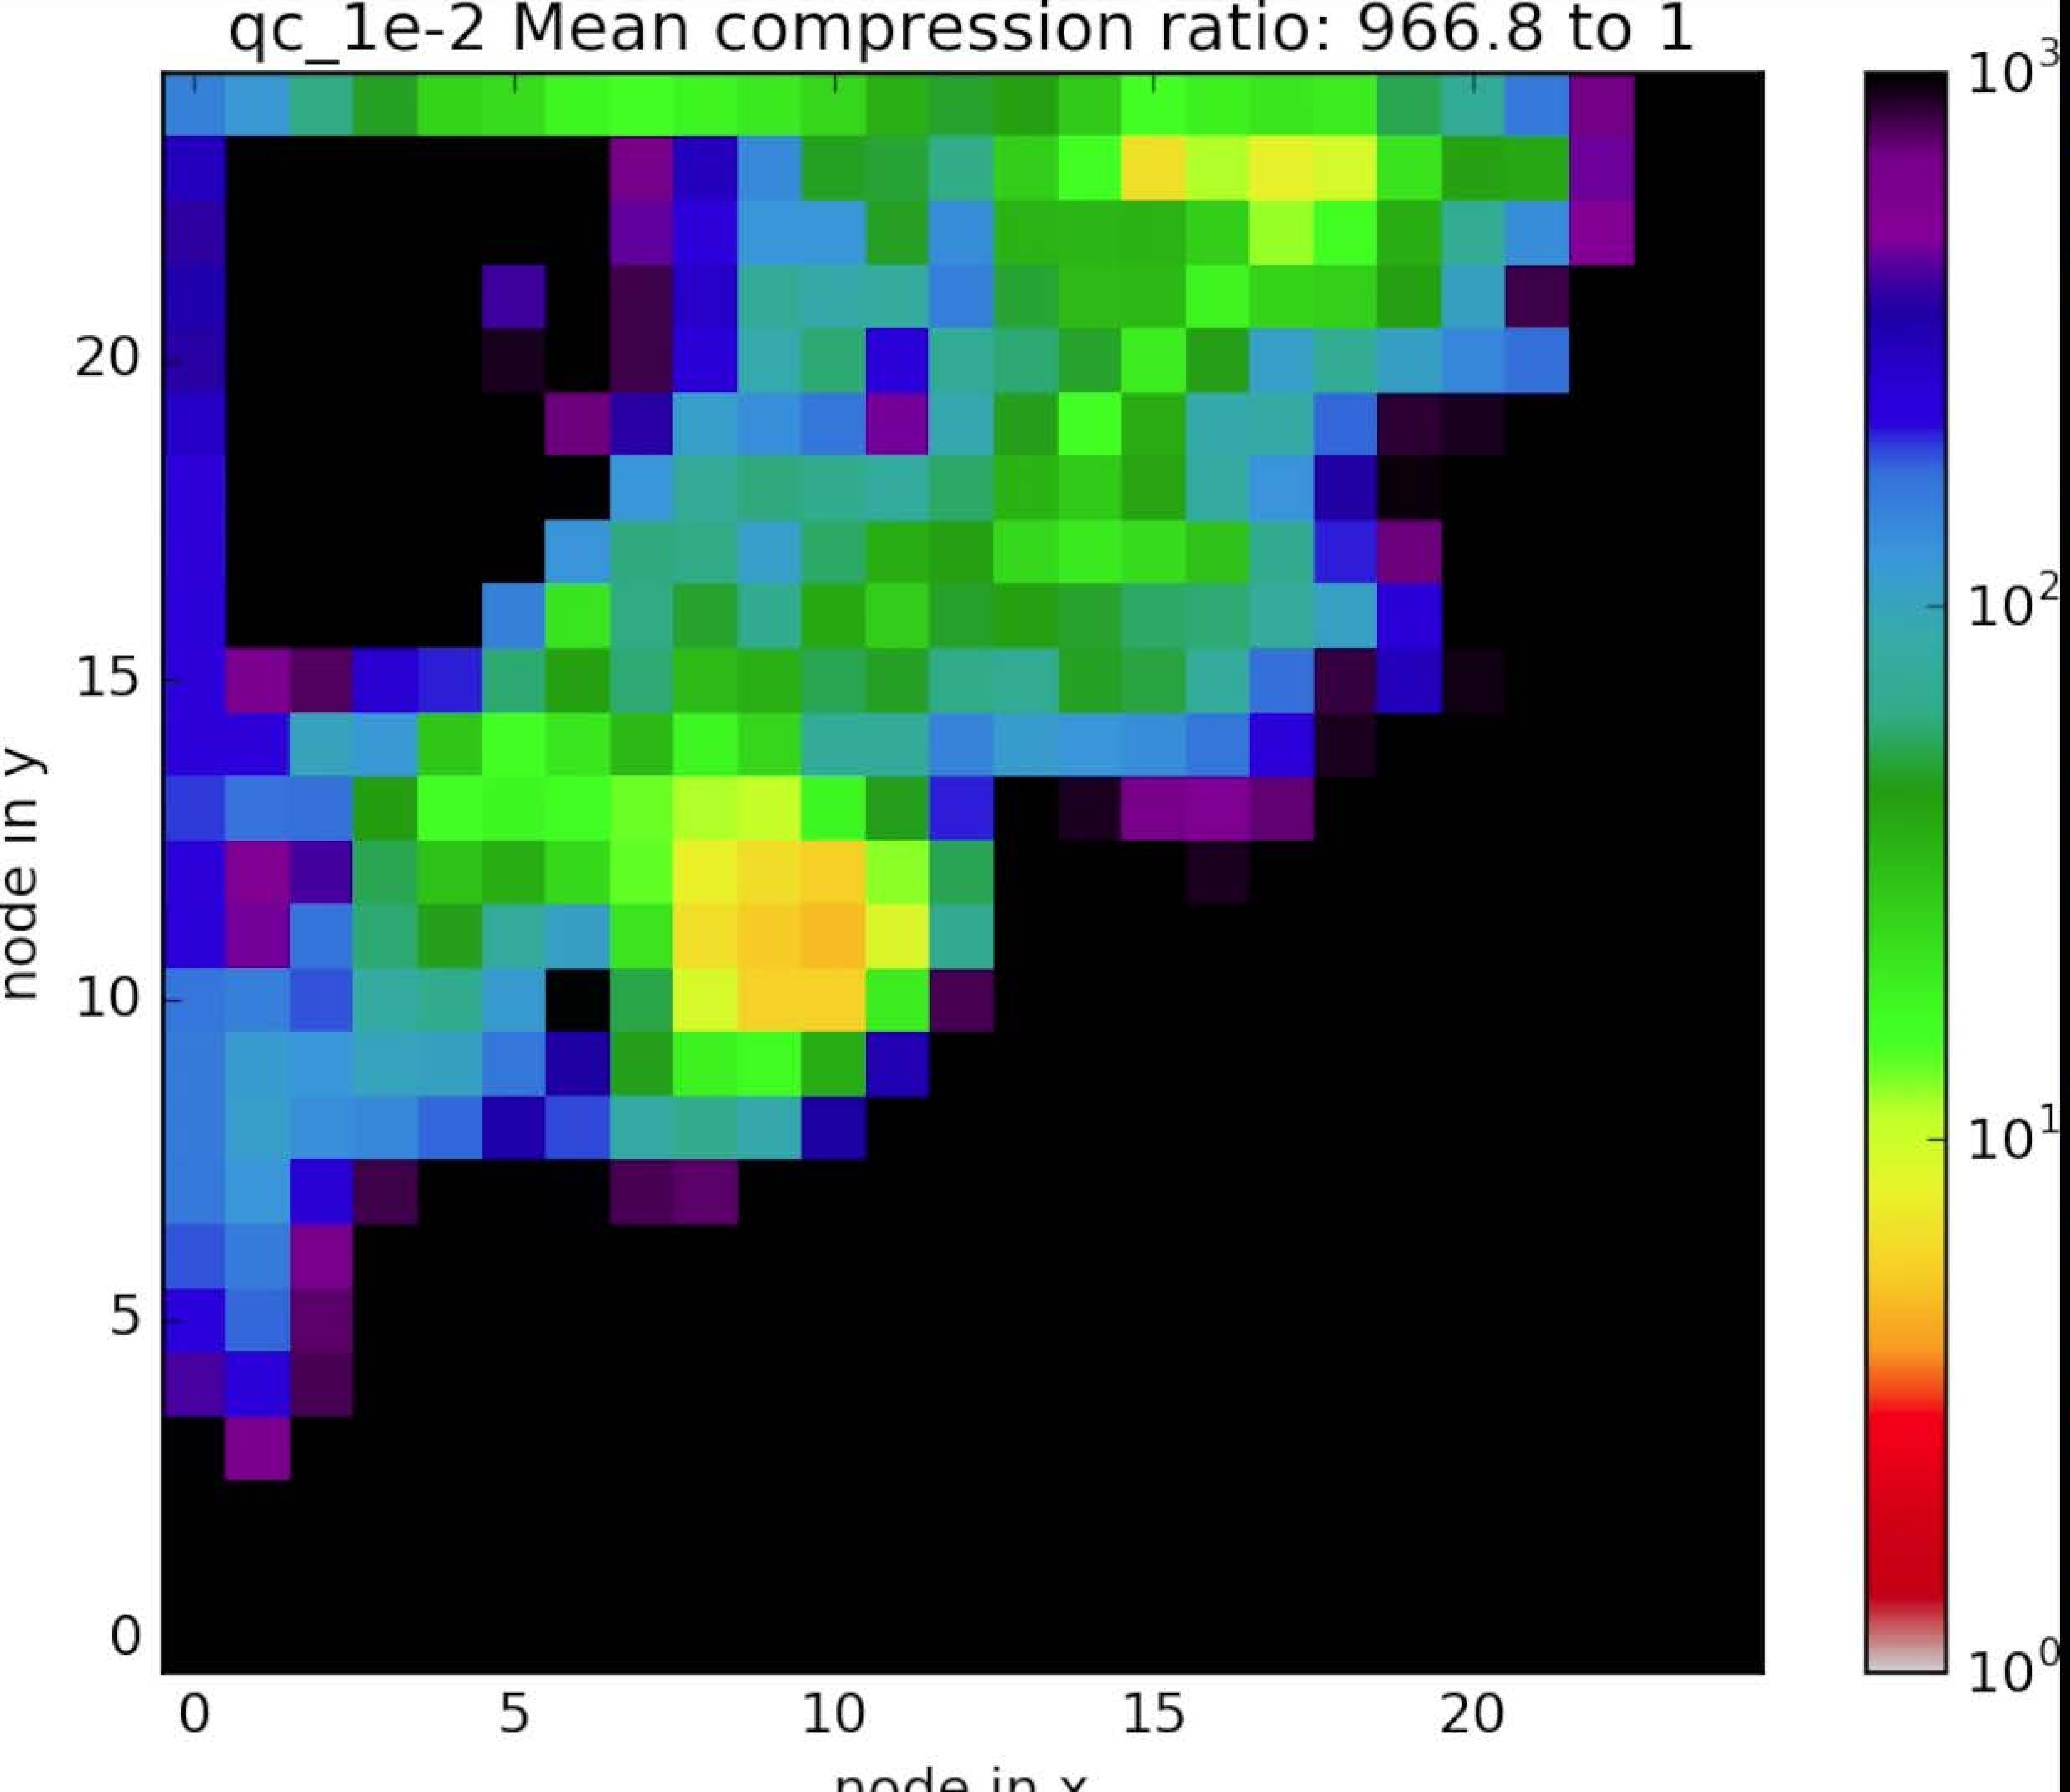

Cloud mixing ratio ranges from 0.0 to 20.0 g/kg

Accuracy=0.01 g/kg

Domain-wide compression ratio 970:1!

However much of the domain does not contain cloud (lots of zeroes)

(note log scale)

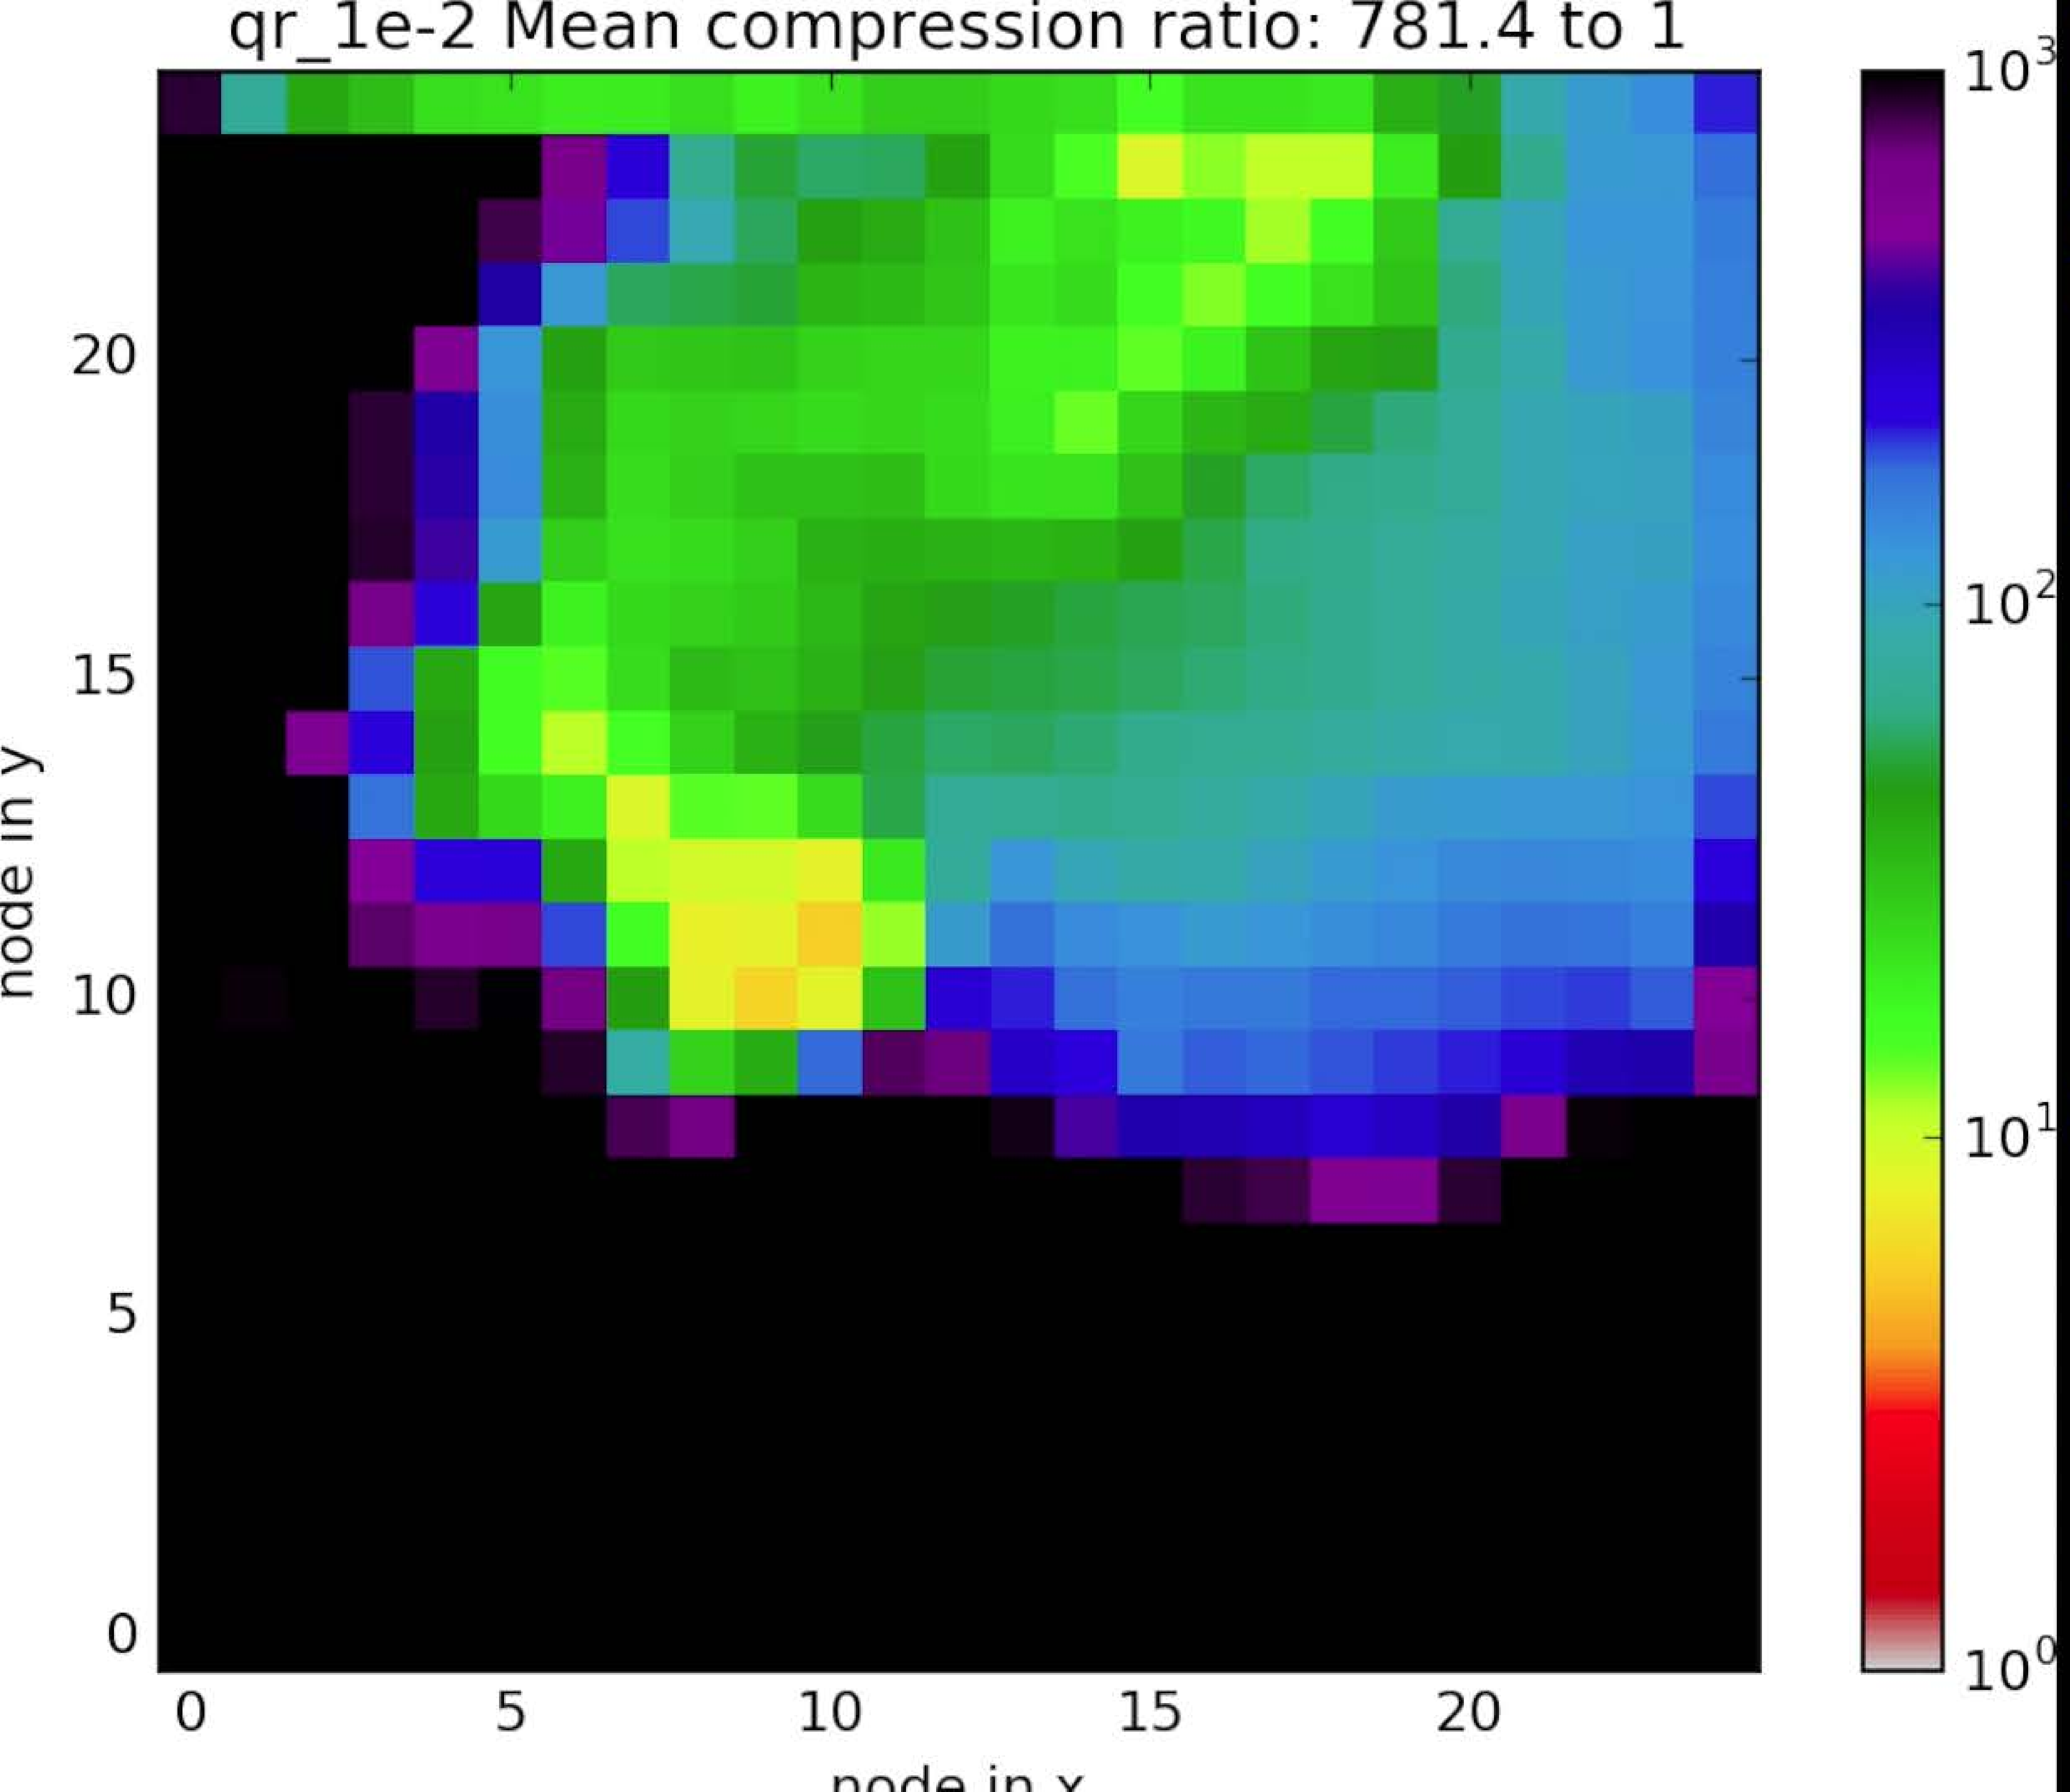

Rain mixing ratio ranges from 0.0 to 20.0 g/kg

Accuracy=0.01 g/kg

Domain-wide compression ratio 780:1!

(note log scale)

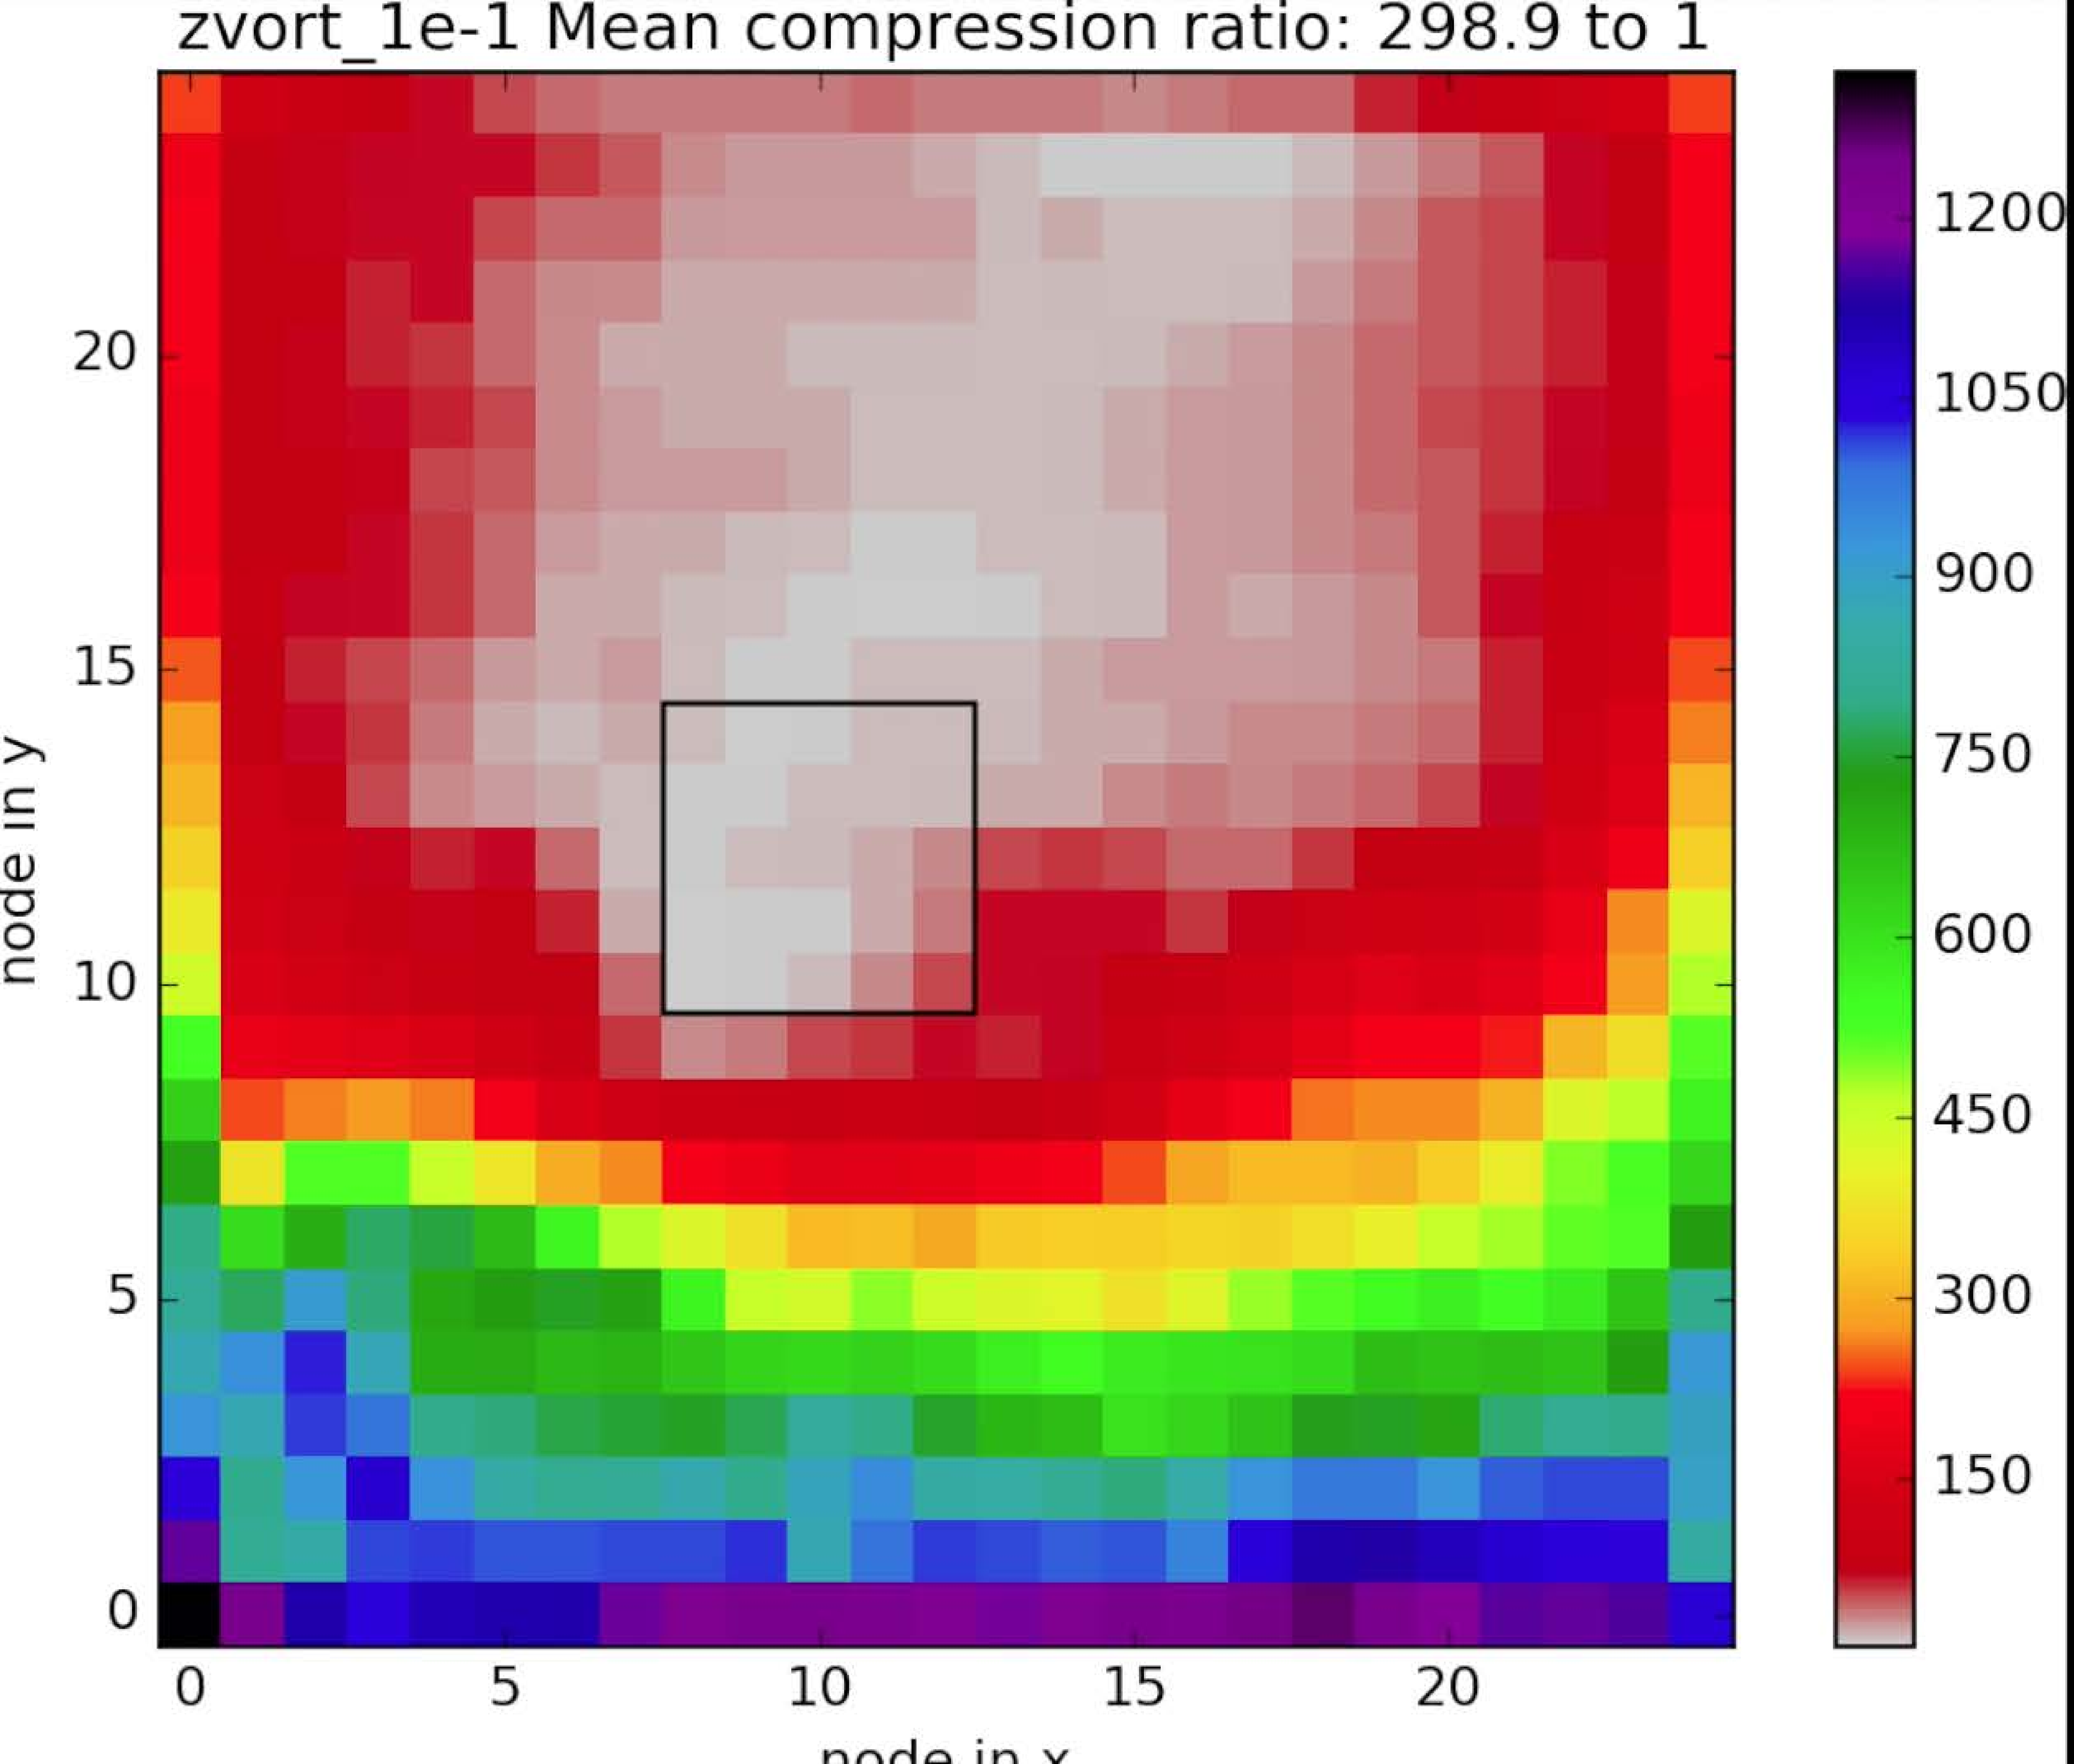

Vertical vorticity ranges from -2 to 4 s^{-1}

Accuracy=0.1 s^{-1}

Domain-wide compression ratio 300:1

This variable exhibits large variation in space

Let's focus on where the action is...

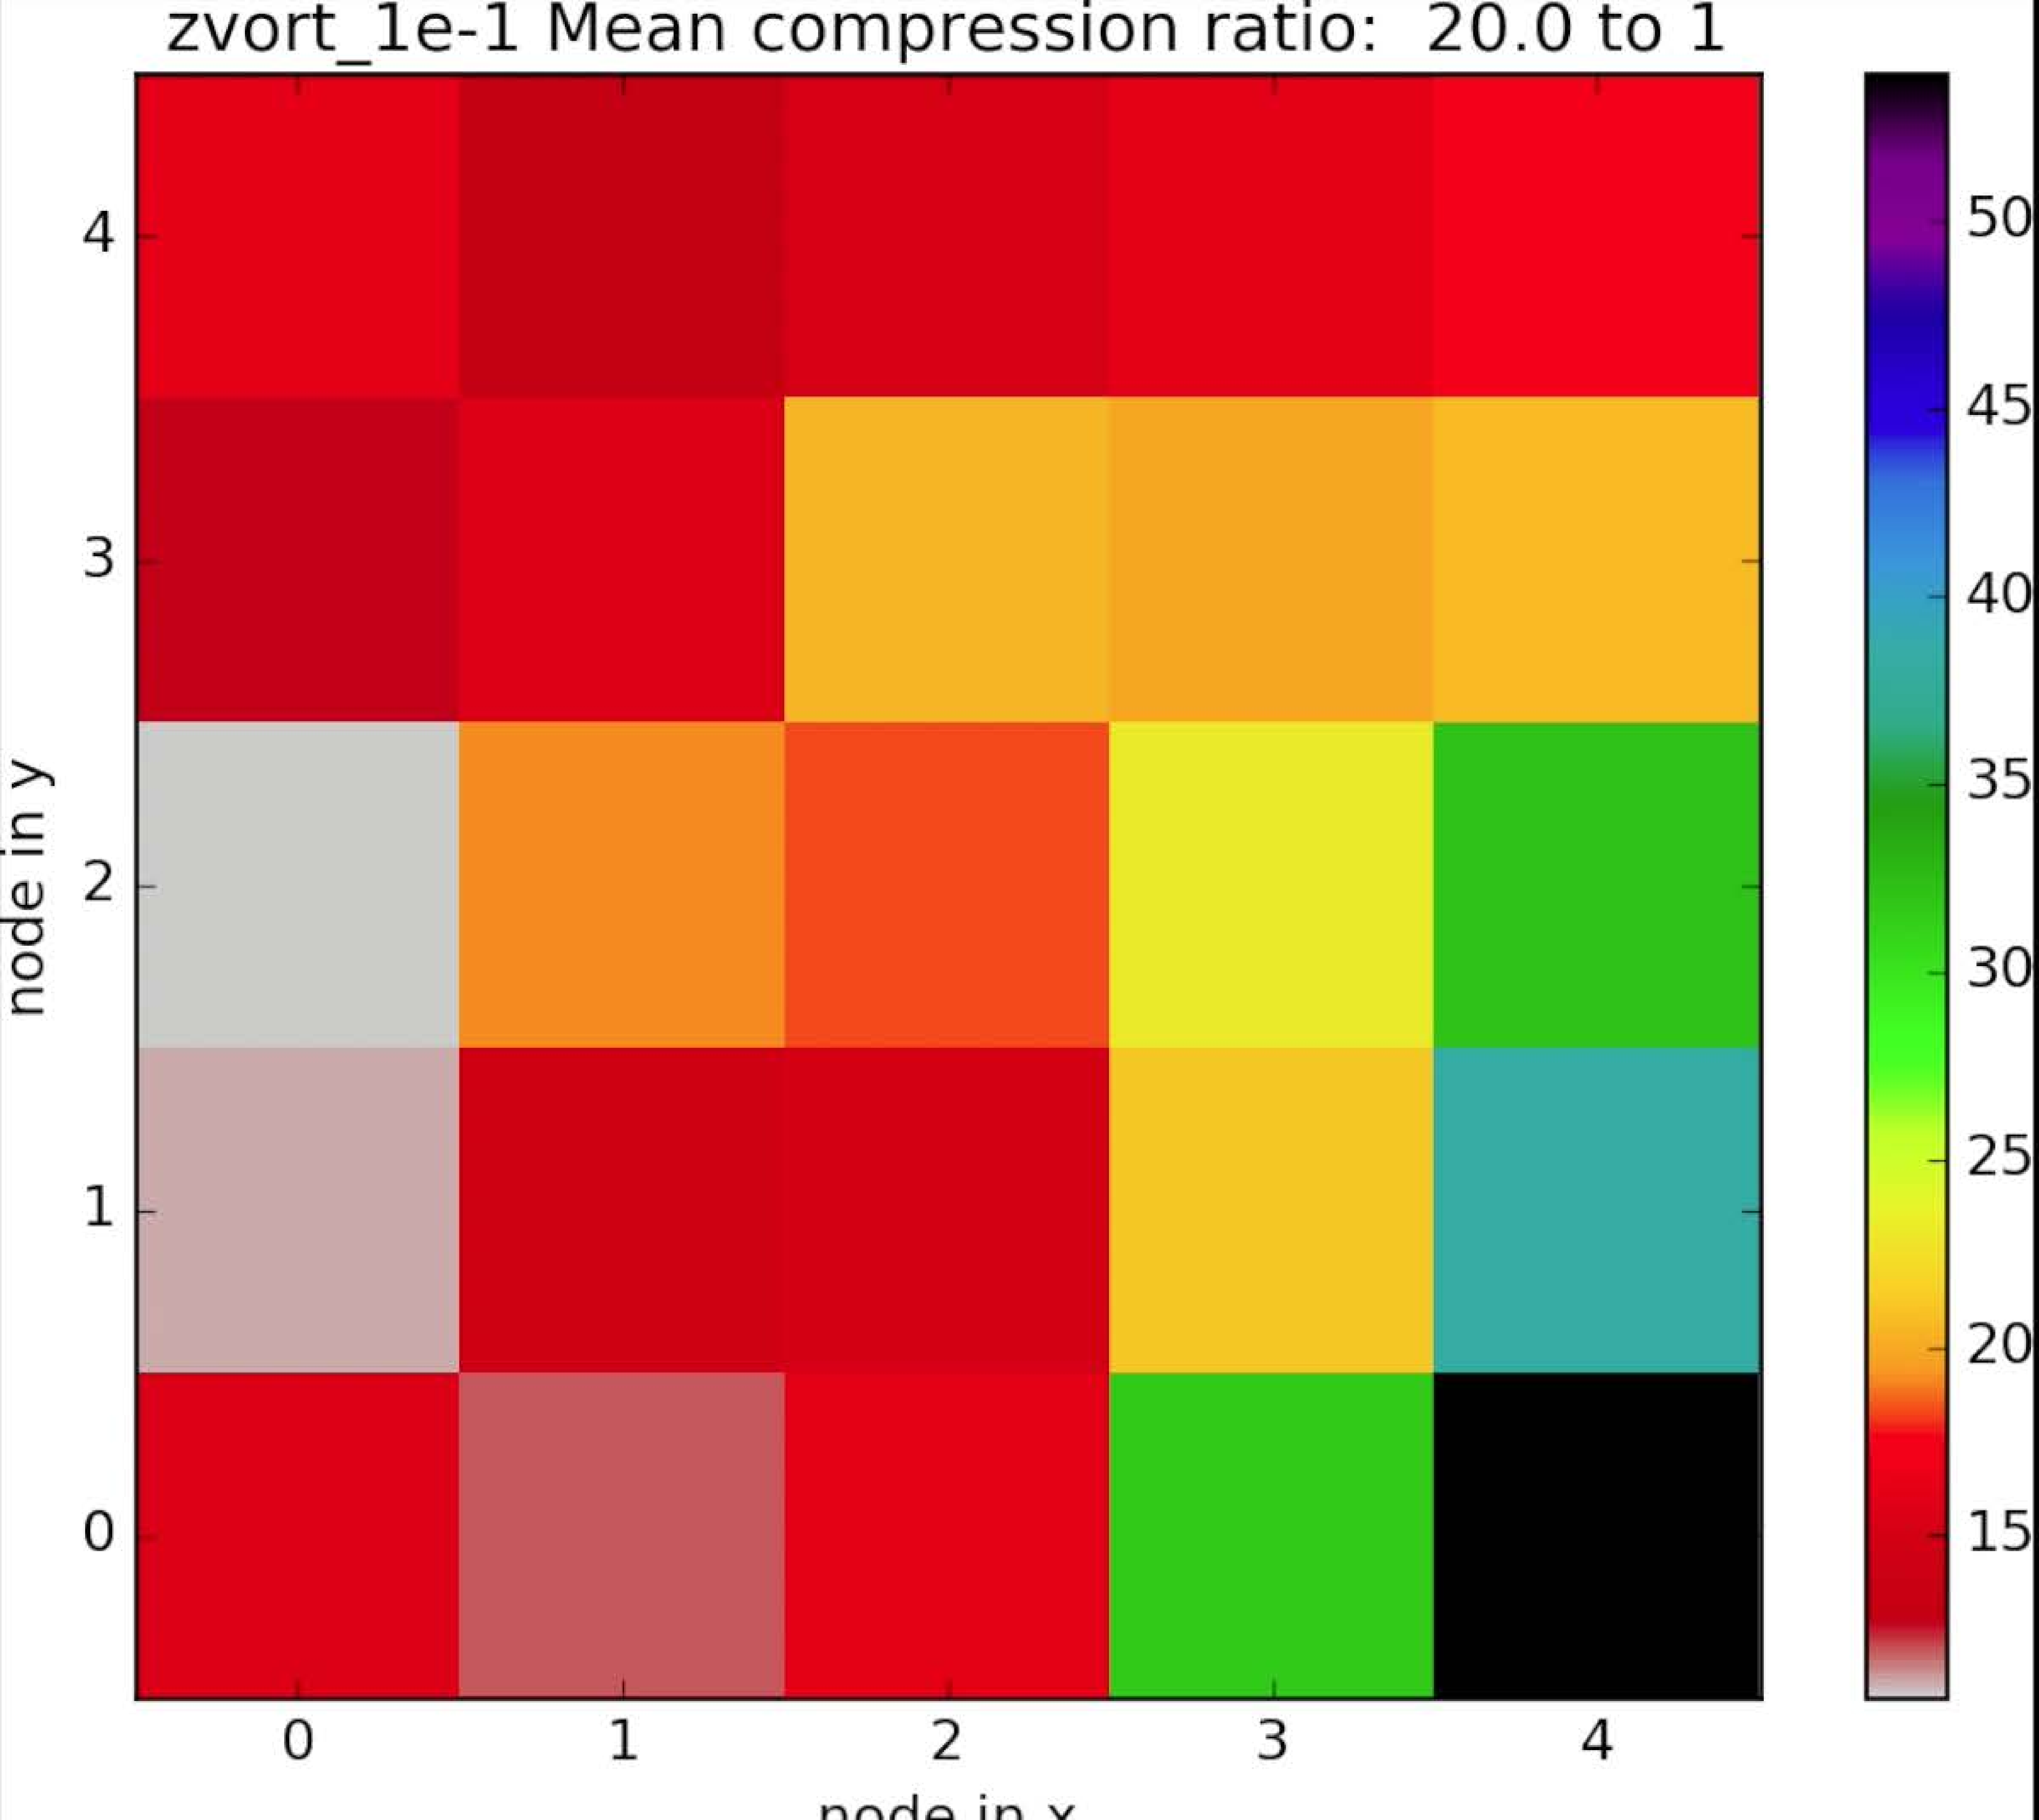

Vertical vorticity ranges from -2 to 4 s^{-1}

Accuracy=0.1 s^{-1}

Local compression ratio 20:1 - in region of most "action"

vertical vorticity
accuracy = 0.1 s^{-1}

vertical vorticity
accuracy = 1.0 s^{-1}

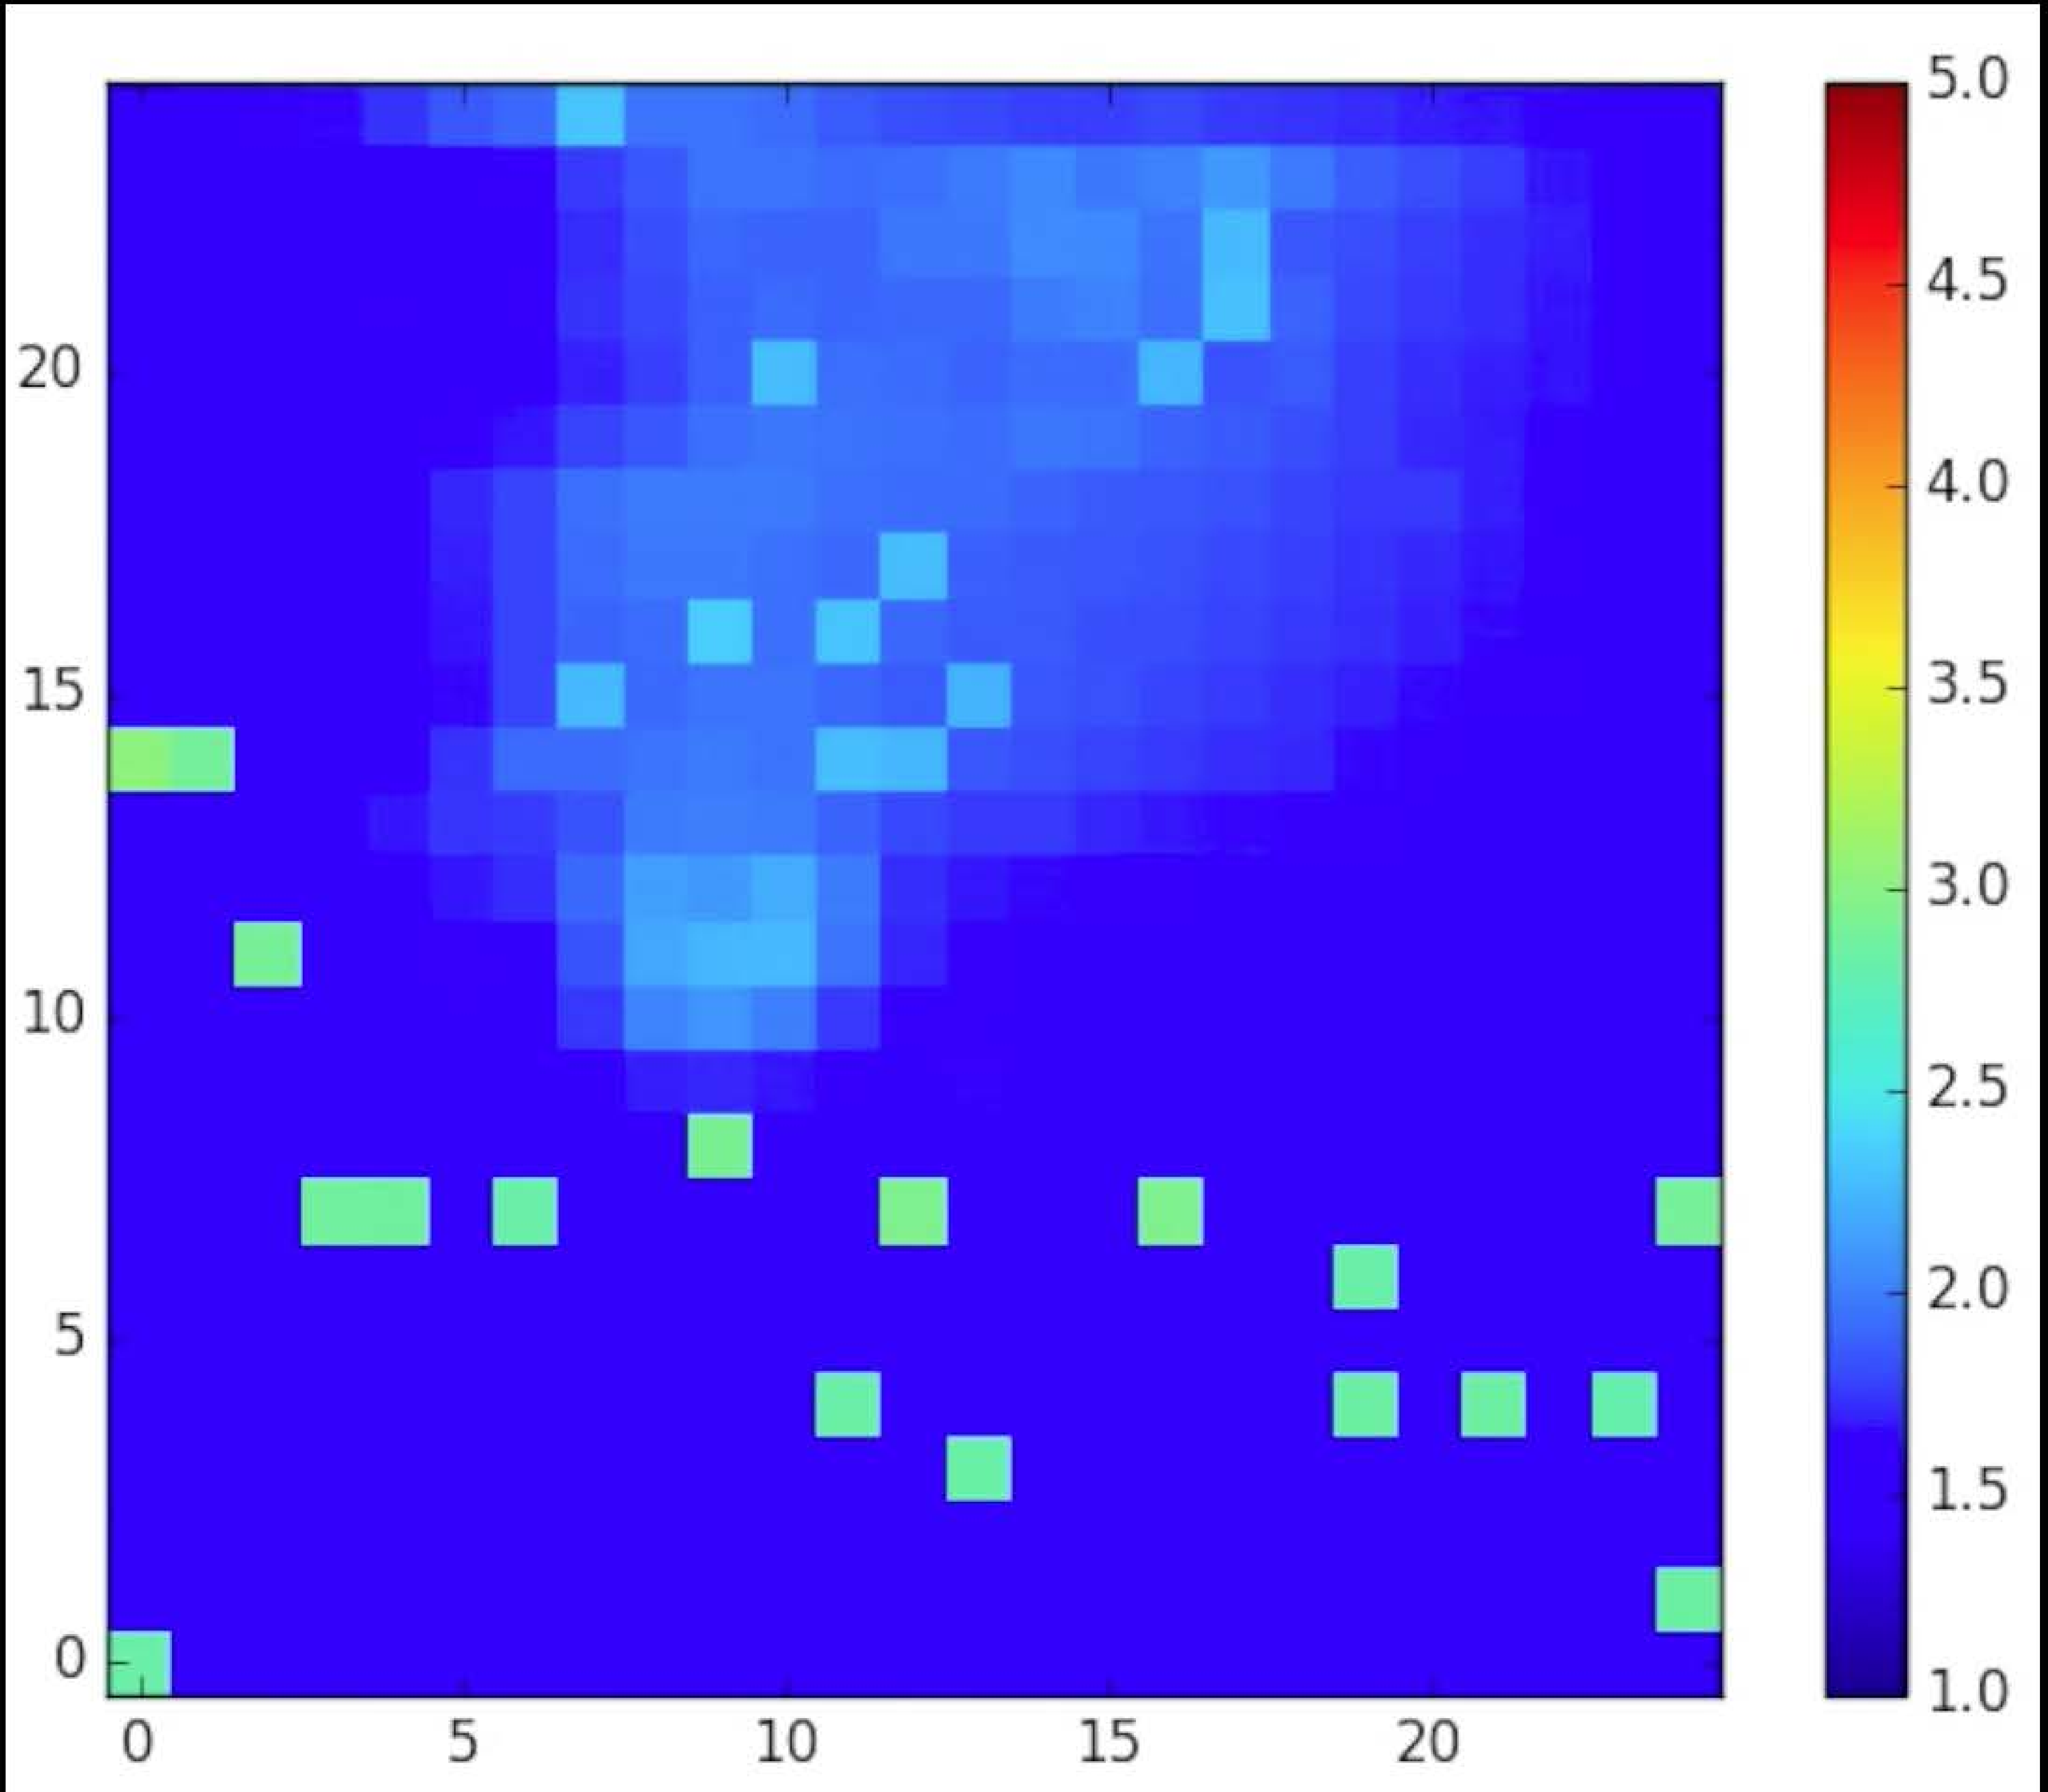

Time (s) to write several ZFP-encoded 3D datasets to memory

Large amount of unexplained variation in write time

Some nodes take 10x longer to buffer data than fastest node

Parting thoughts

- It is now possible to run full-physics simulations at resolutions approaching that of previous chamber studies
- Utilizing ZFP compression has opened the door for us to do more analysis, and helps reduce the data load on Blue Waters
- Features such as the SVC, VVS, SVS, the “parade of vortices” along the FFDB are the focus of current analysis. *Do these things exist in nature?*
- Future work will include sensitivity studies within the current environment, as well as other environments observed adjacent to observed long-path EF5 tornadoes

Thank you!

See `http://orf.media` for more

Questions?