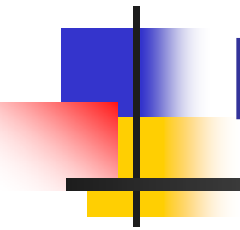

Expanding Vapor's Data Handling Capabilities

Victor Snyder

NCAR SIParCS

August 10, 2007

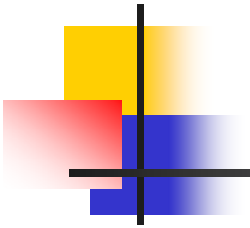

Introduction

- Handle data from Weather Research and Forecasting (WRF) model
  - Non-uniform, non-rectangular grid
- Direct output to Vapor's format
  - Parallel wavelet transformation

Evaluating Vapor and WRF

- Weather data rarely seen in 3D
- Preliminary tests: WRF data looks interesting in 3D

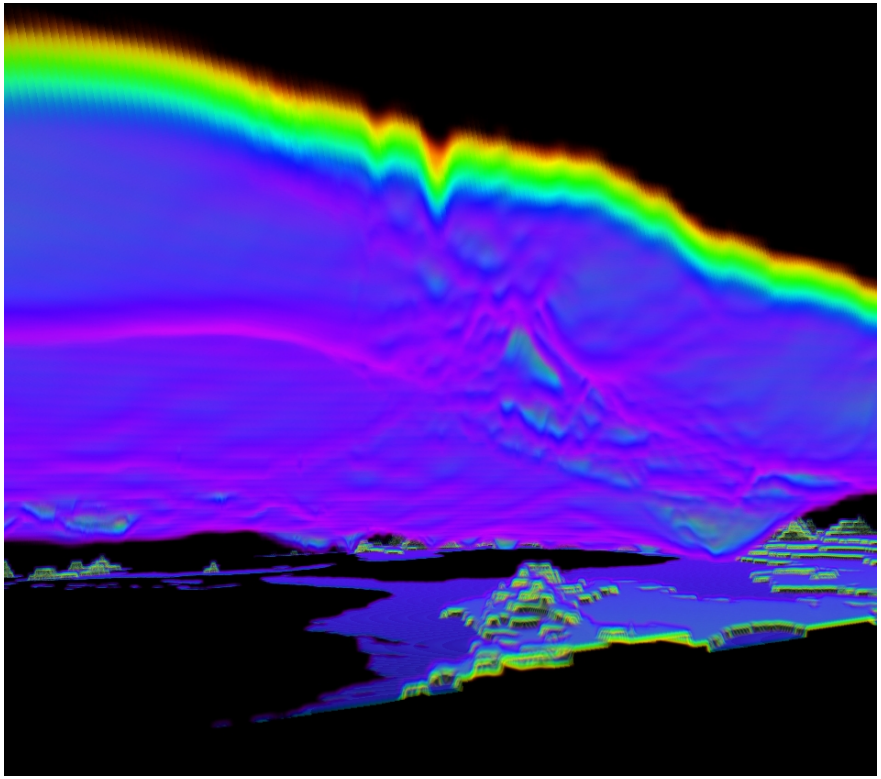

WRF User's Requirements

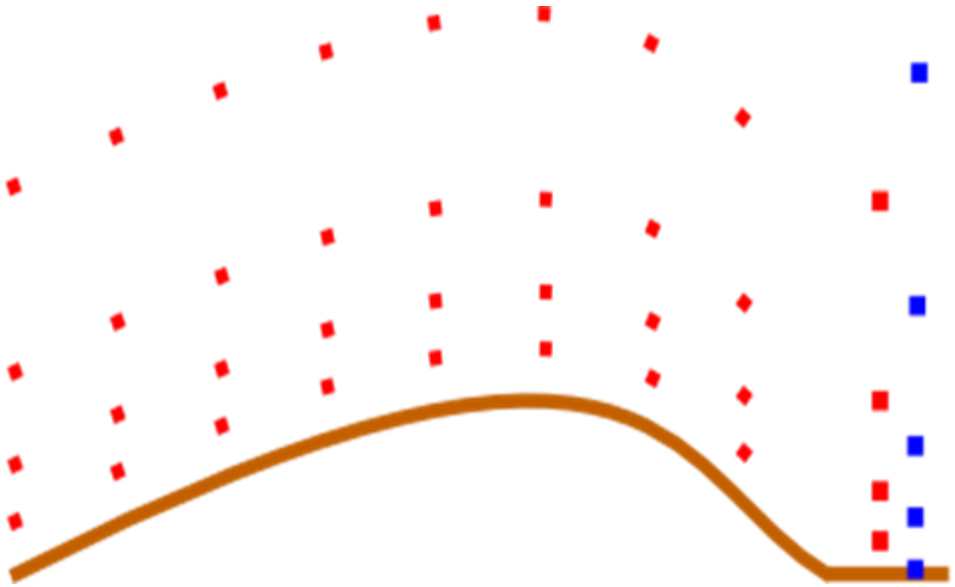

Vapor should:

- Handle odd grids
- Deal with invalid data points
- Add certain derived variables
- Handle time steps

And do this all quickly

WRF Data Conversion

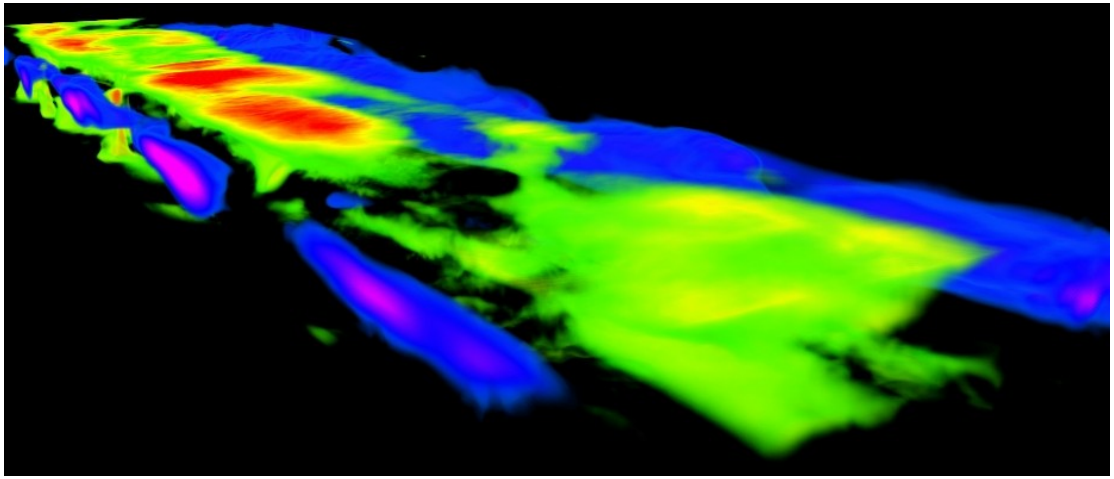

- Correctly handles staggered grids
- Efficiently adds derived quantities
- Uses WRF-specific metadata to automate process

Vapor and Large Data Sets

- Wavelet transforms
  - “Smooth” and “detail” parts
- Data conversion
  - Time consuming
  - Redundant
- Goal: simulations output Vapor’s format

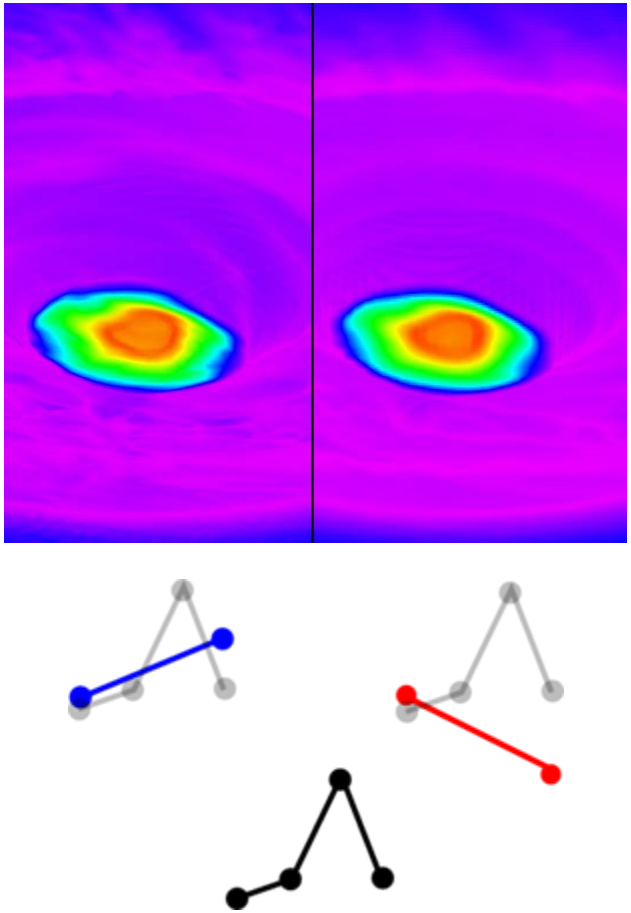

Parallel Vapor Output

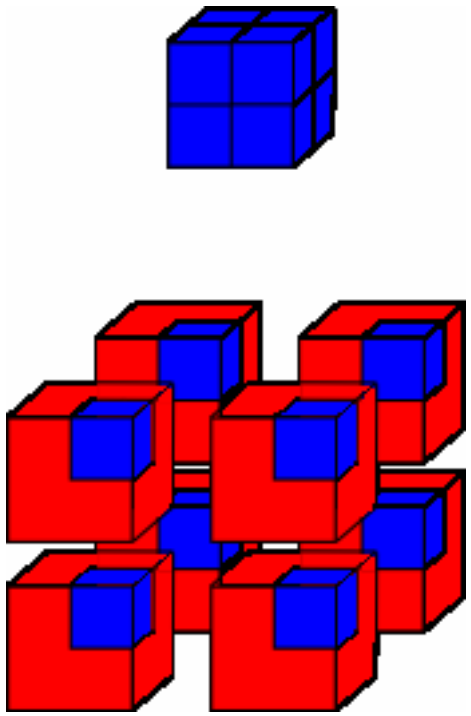

- Communication
  - Move data points into blocks
- Transformation
  - Readily parallelized
- Output
  - Parallel netCDF
  - netCDF 4
  - HDF5

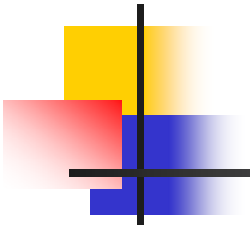

Initial Tests

- Simplified situation
  - Data already arranged properly
  - Haar wavelets
- Code for transformation of block on single node
  - 8 processors, 2^{18} points: ~5x speed-up

Conclusion

- Vapor's improved data handling
  - Vapor useful to WRF users
  - Direct output to Vapor's format is viable