Access to European High Performance Computing – A PRACE User Perspective

Åke Nordlund

Niels Bohr Institute
University of Copenhagen

Overview

- Context
  - User profile – experiences
- PRACE Projects
  - Solar research
  - Other grant examples: The supernova project
- PRACE-related e-Infrastructure Requirements
  - Tier-0 / Tier-1 access
  - Data ***analysis***, ***archival***, and peer ***access***

User Profile

- 40 years HPC experience
  - 1970's: UNIVAC, CDC
  - 1980's: IBM, Convex, …
  - 1990's: Connection Machine (CM-200)
  - 2000's: NASA/Ames, Stuttgart, Jülich (JUGENE)

- Danish Center for Scientific Computing
  - DCSC-KU: Intel-cluster, GPU-cluster (200 GPUs)
  - ScaleMP Virtual Shared Memory system
    - 'Analysis server'

Solar Modeling

The magnetic field in the solar corona is controlled by the magnetic field at the surface of the Sun

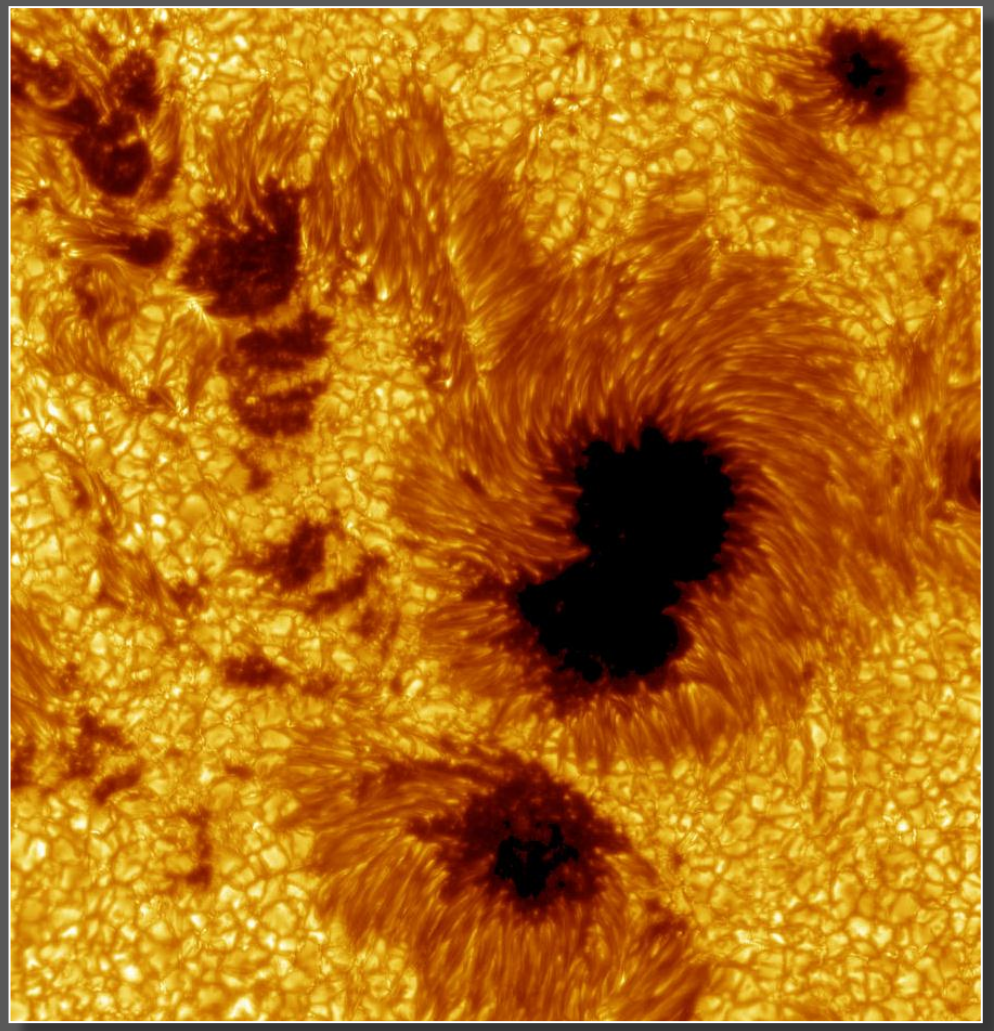

A Nearly Scale Free Spectrum!

Doppler Image of the Sun

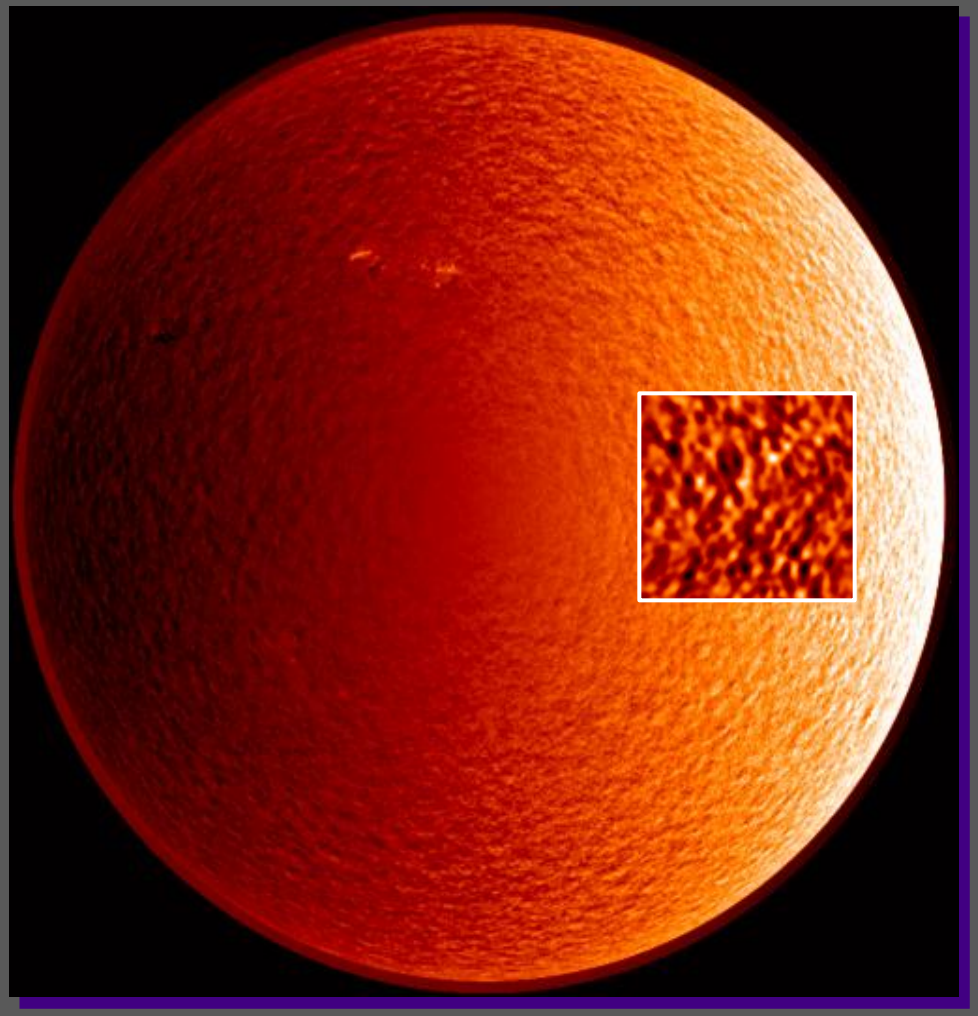

Solar horizontal velocity (observed)

Scales differ by factor 2 – which is which?

400 Mm

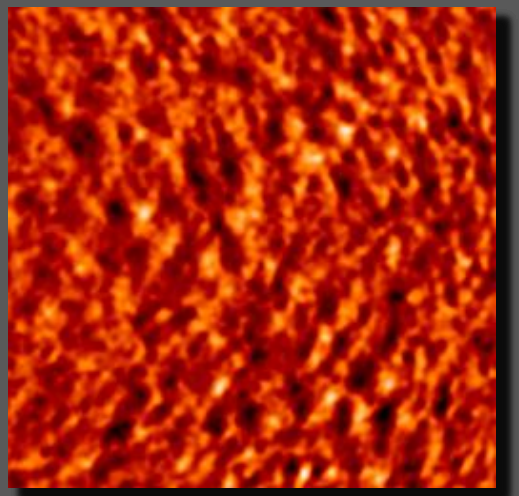

100 Mm

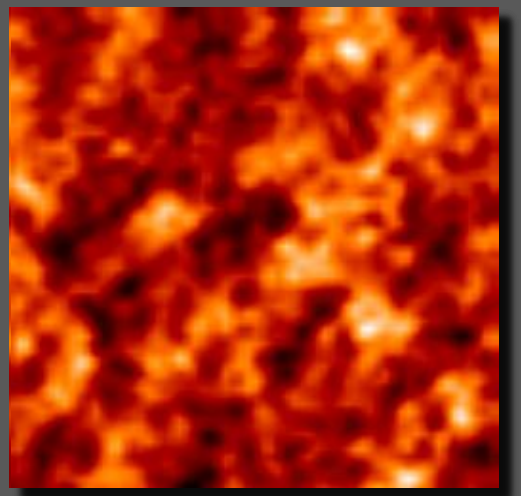

50 Mm

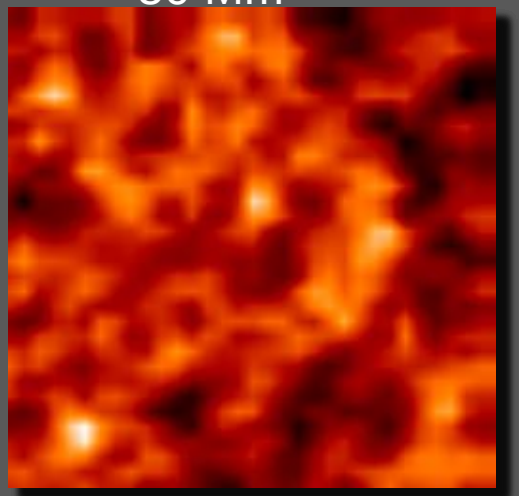

200 Mm

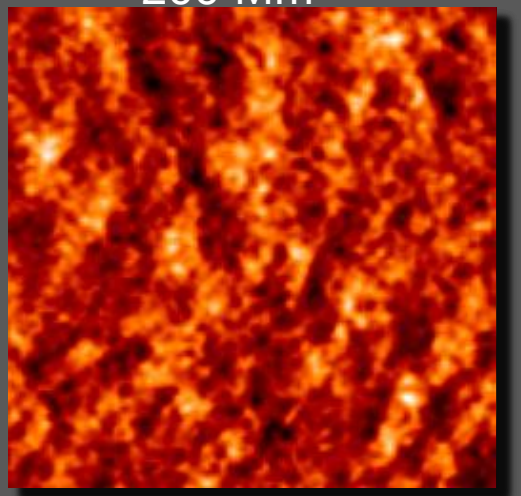

Solar horizontal velocity (model)

Scales differ by factor 2 – which is which?

12 Mm

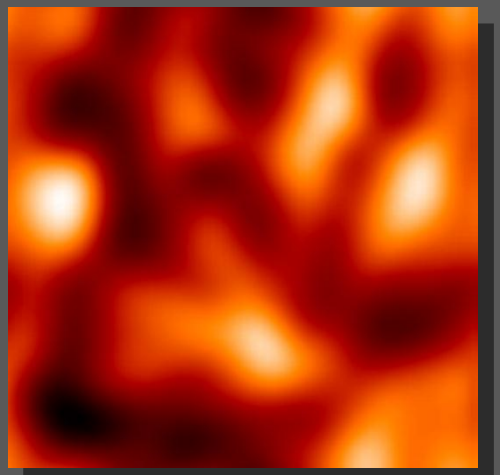

24 Mm

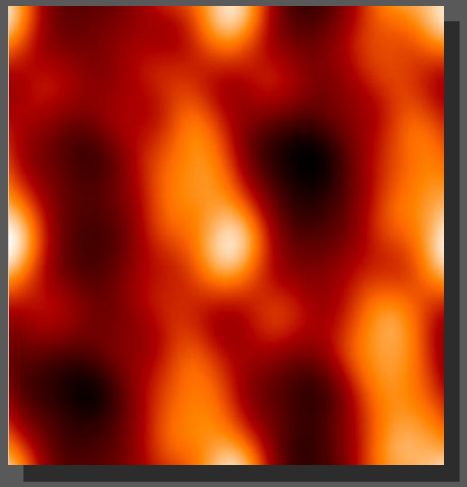

3 Mm

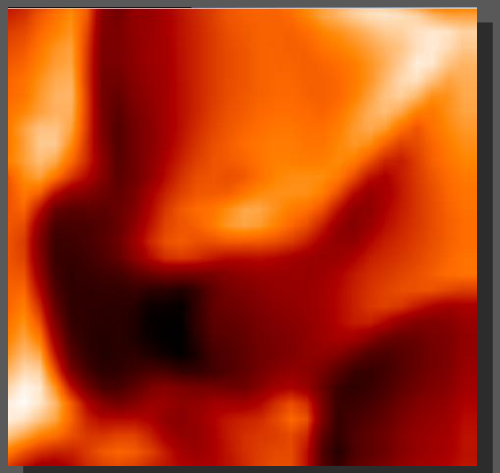

6 Mm

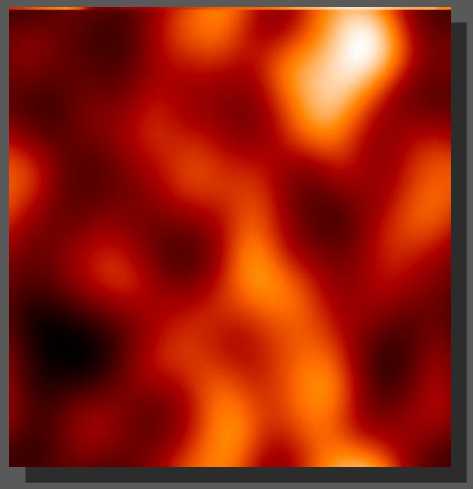

Solar velocity spectrum

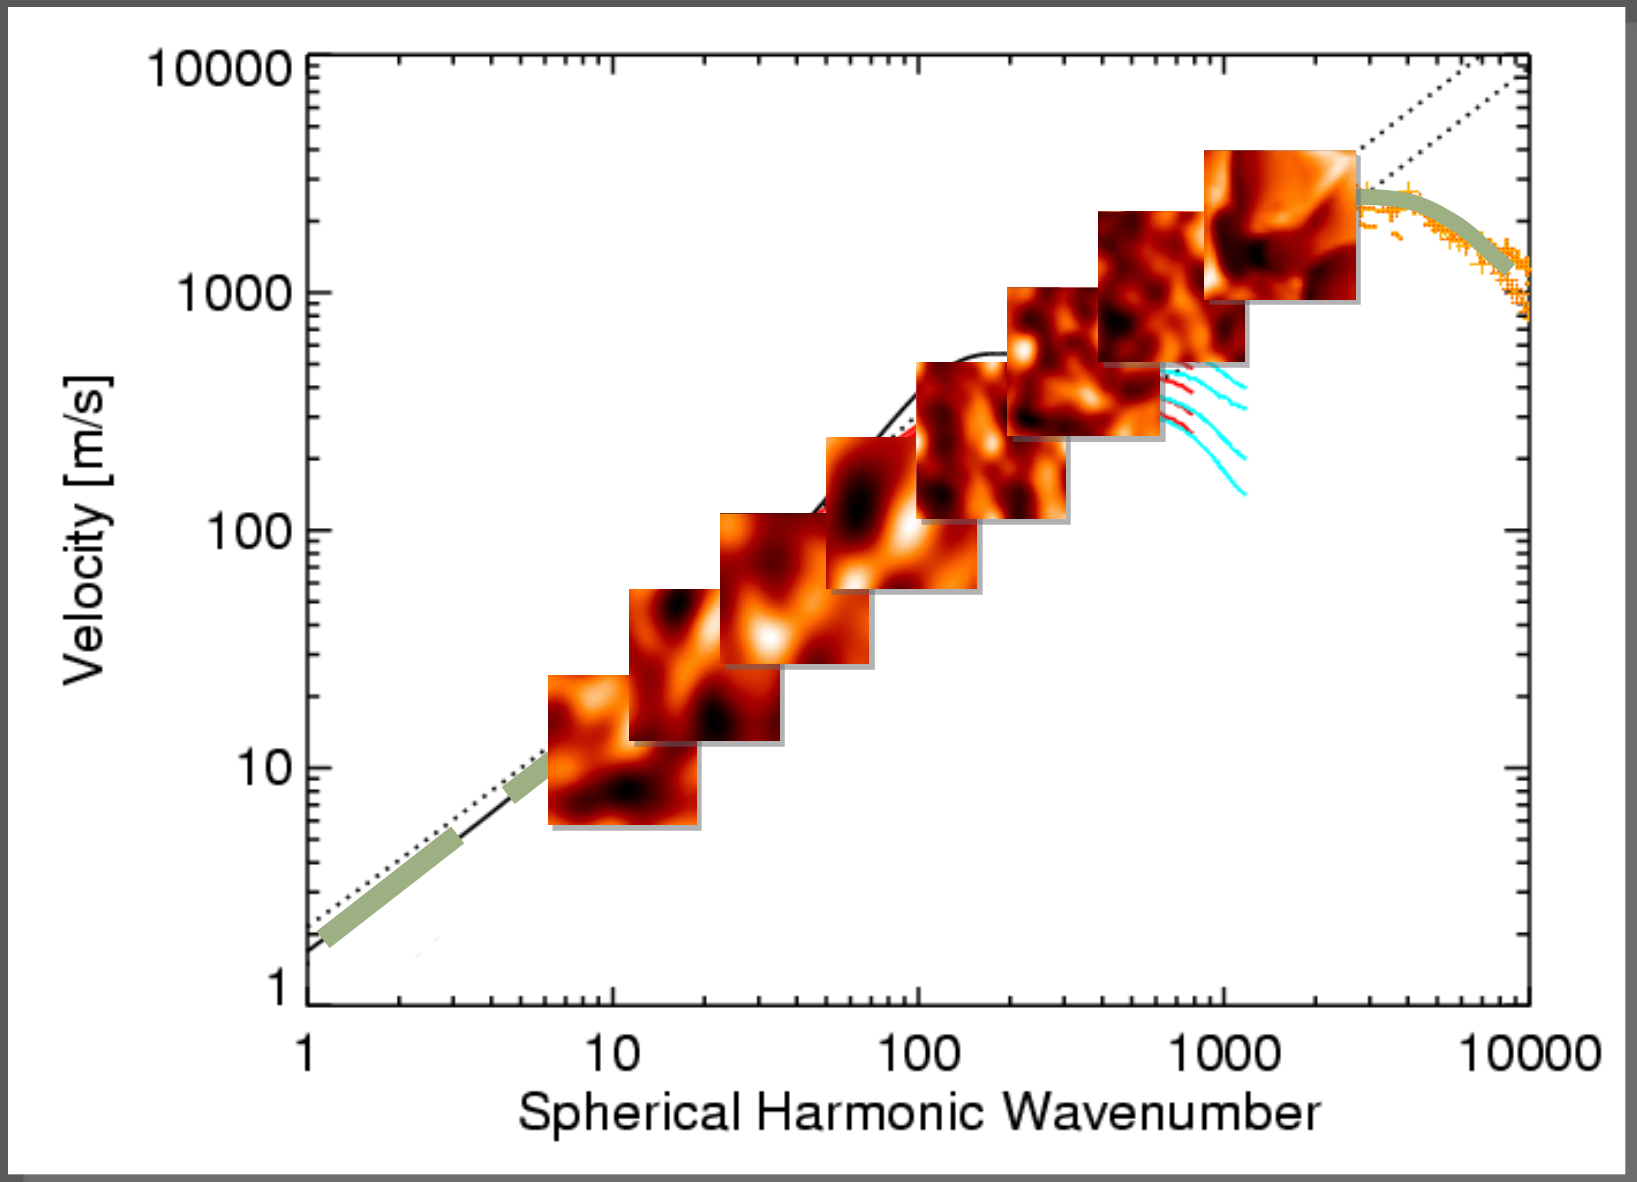

Solar Active Region 3-D Simulations

- **'Quiet' Sun**
  - Zero mean field with $\langle B^2 \rangle^{1/2} \sim 50 - 150$ G
- **Plage Region**
  - Mean vertical field **B ~ 600 G**
  - 24,000 km size
    grid sizes **down to 6 km**
- **Active Region with Sunspots**
  - Zero mean field with $\langle B^2 \rangle^{1/2} \sim 1.5$ kG
  - 48,000 km size,
    horizontal grid size 24 km

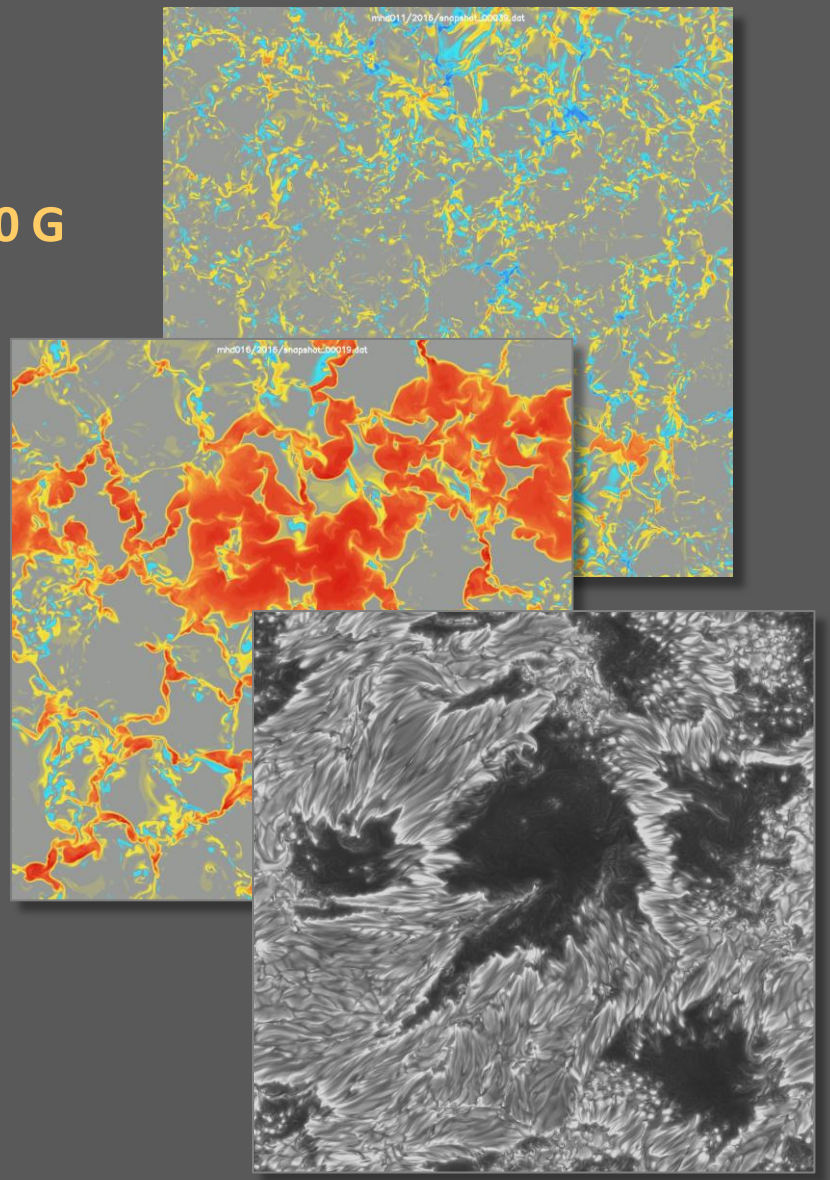

Solar 'Plage' Regions

Line-of-sight velocity

Continued spot evolution (simulations at NASA/Ames by Bob Stein)

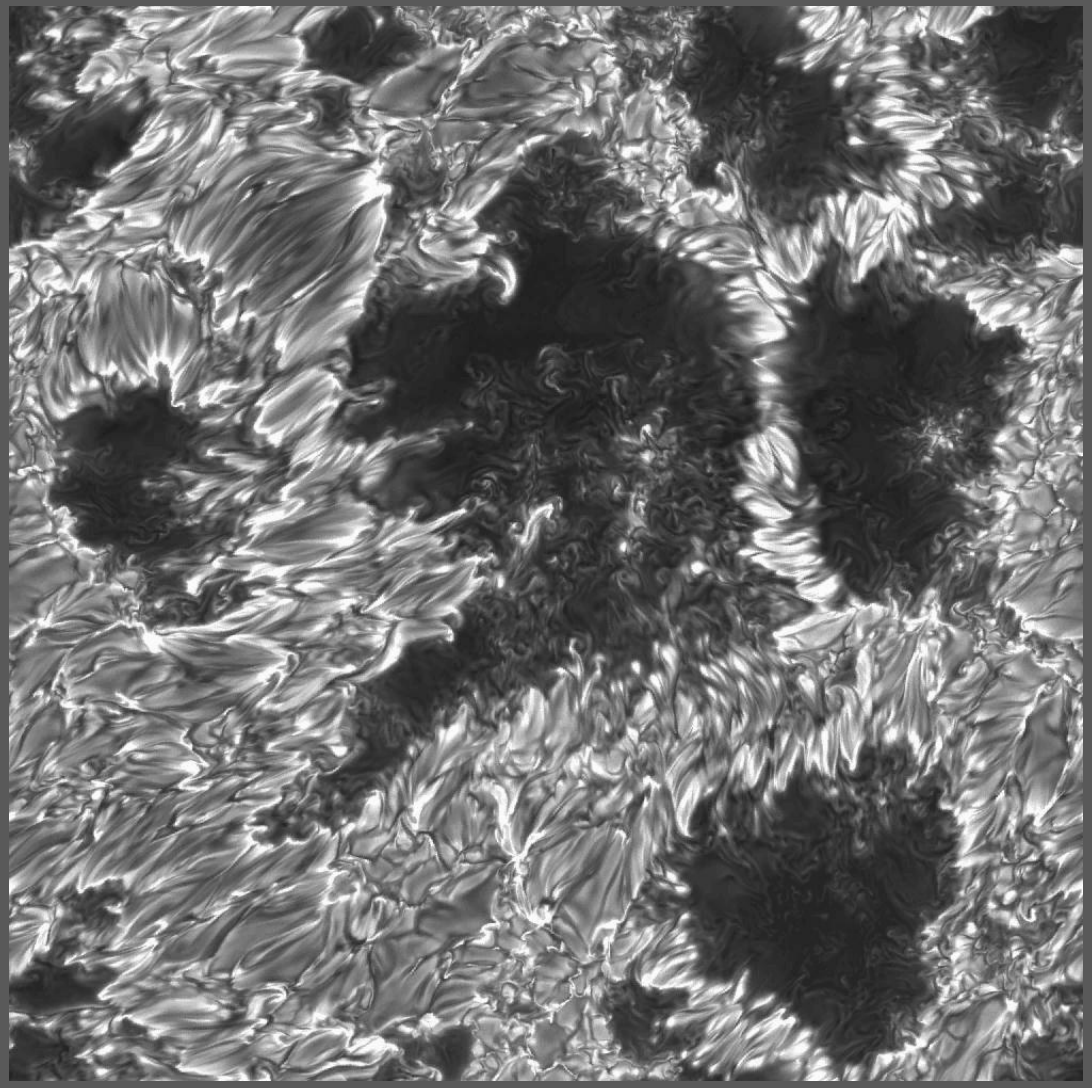

Size: 48^2 x10 Mm
Mesh: 2016^2 x 500

Continued Spot Evolution
Zoom in on the central spot

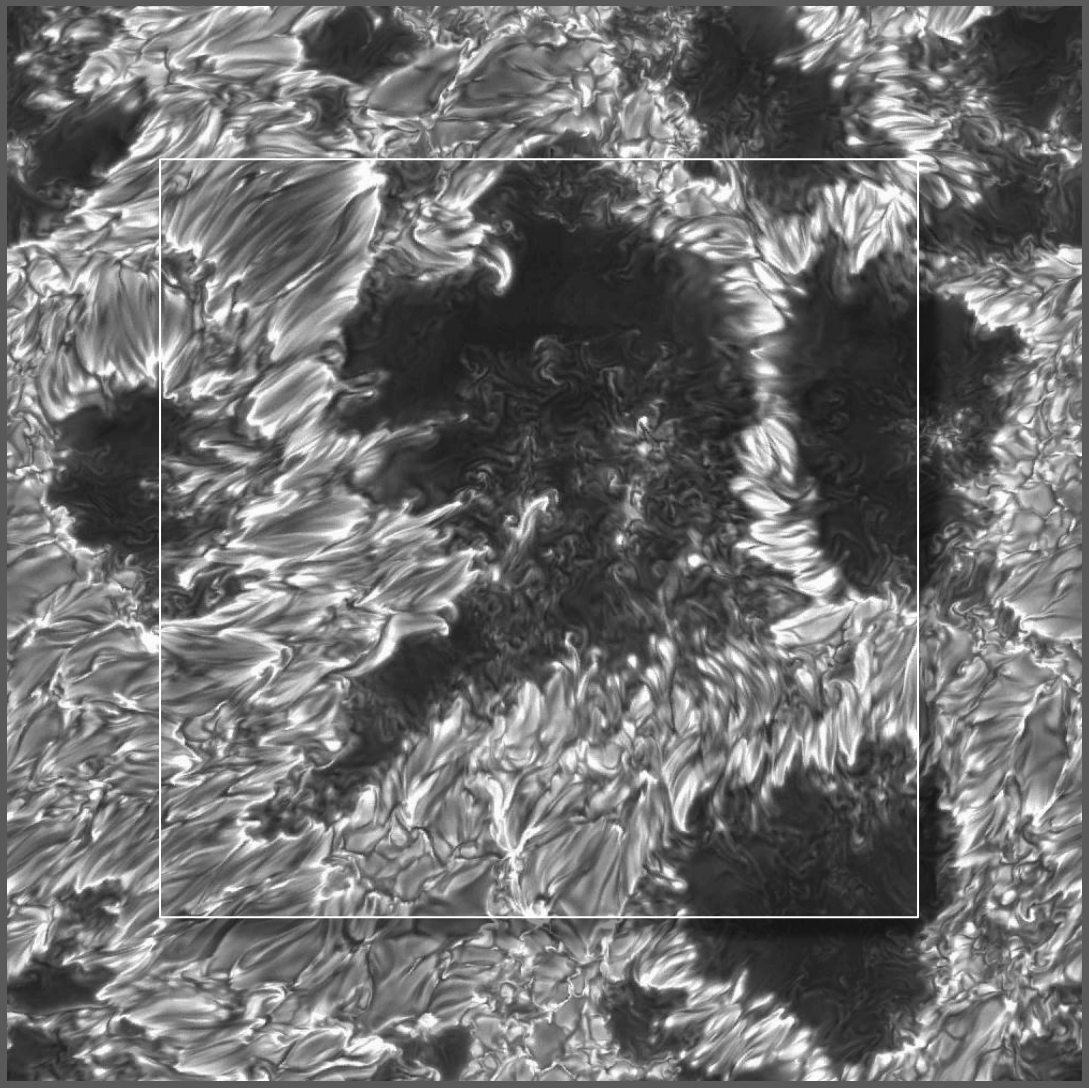

Continued Spot Evolution
Zoom in on the spot at lower right

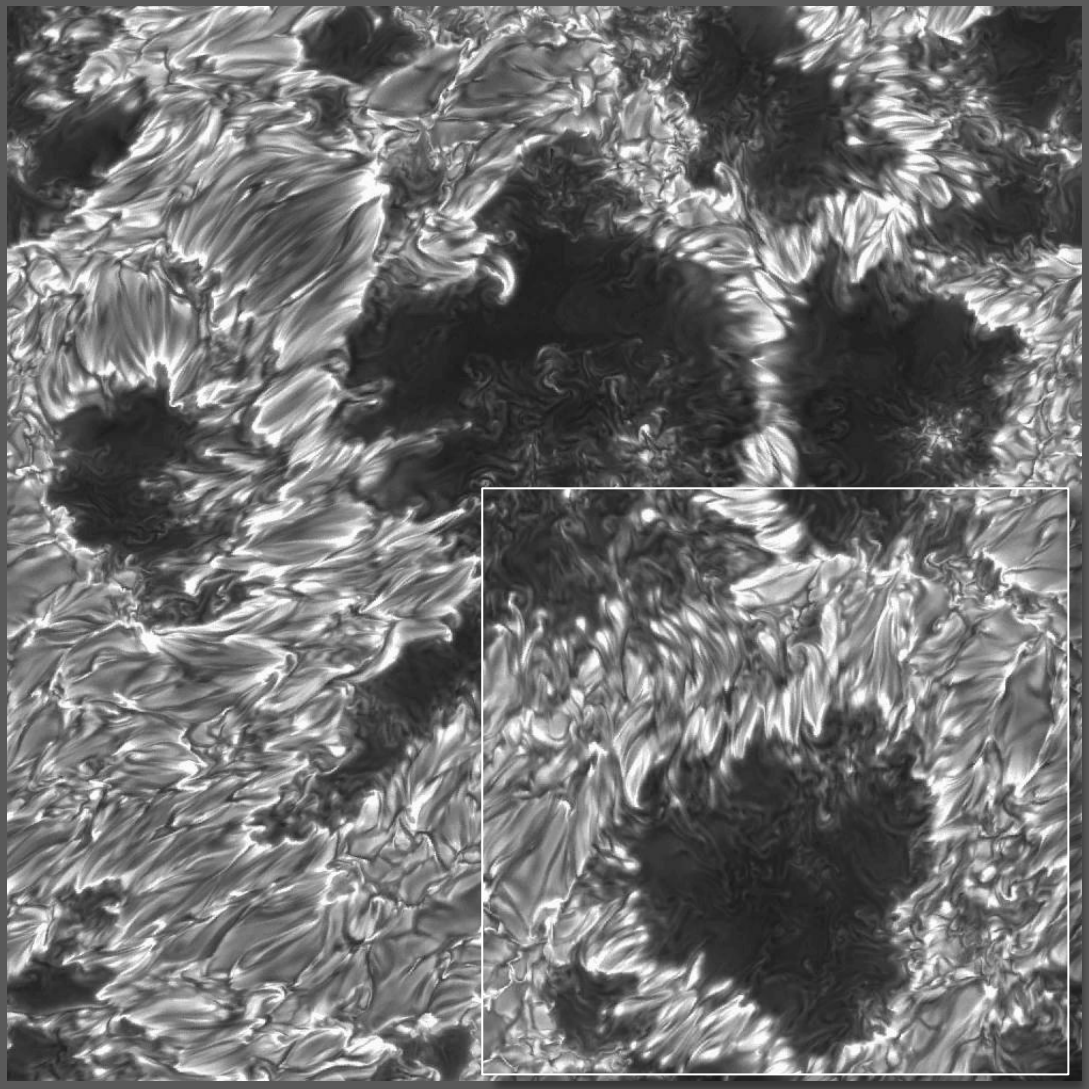

Importance and Impact of Visualization

- Importance:
  - Visualization in 3-D, including time animation is crucial
  - Complements quantitative / statistical analysis
- Impact:
  - High cadence fluid snapshots, each ~0.1-0.2 TB
  - High cadence sub-sampled particle snapshots, each ~1 TB
    - Restart snapshots: 30 TB
- Requirements:
  - High data I/O bandwidth!!
  - Wavelet compression from server to user

3-D (NCAR/Vapor) visualizations, illustrating the process

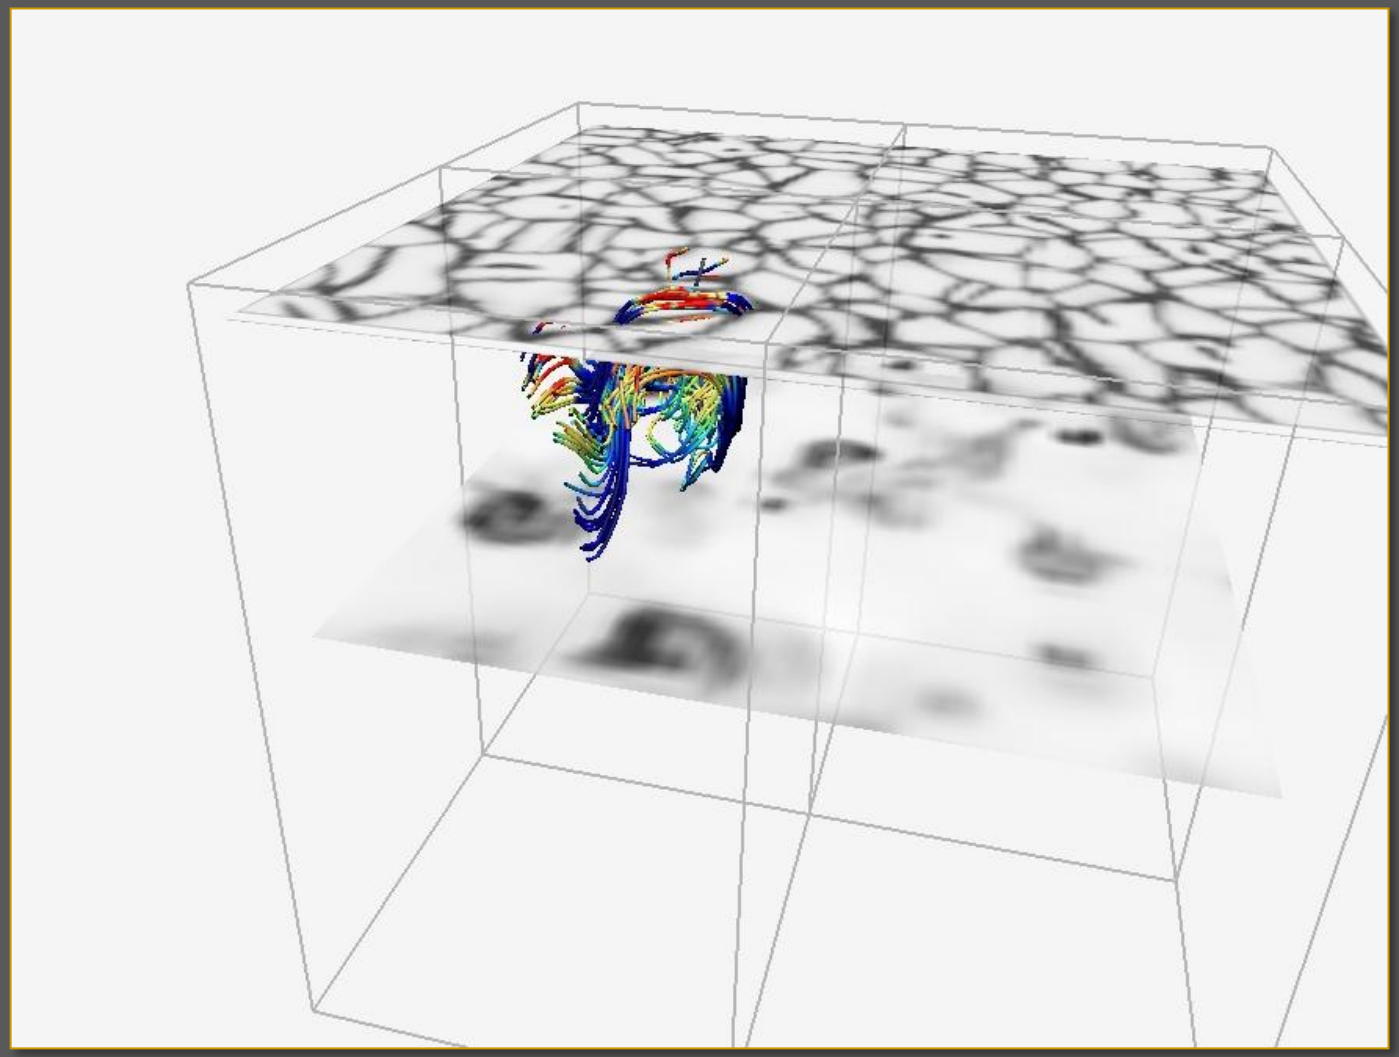

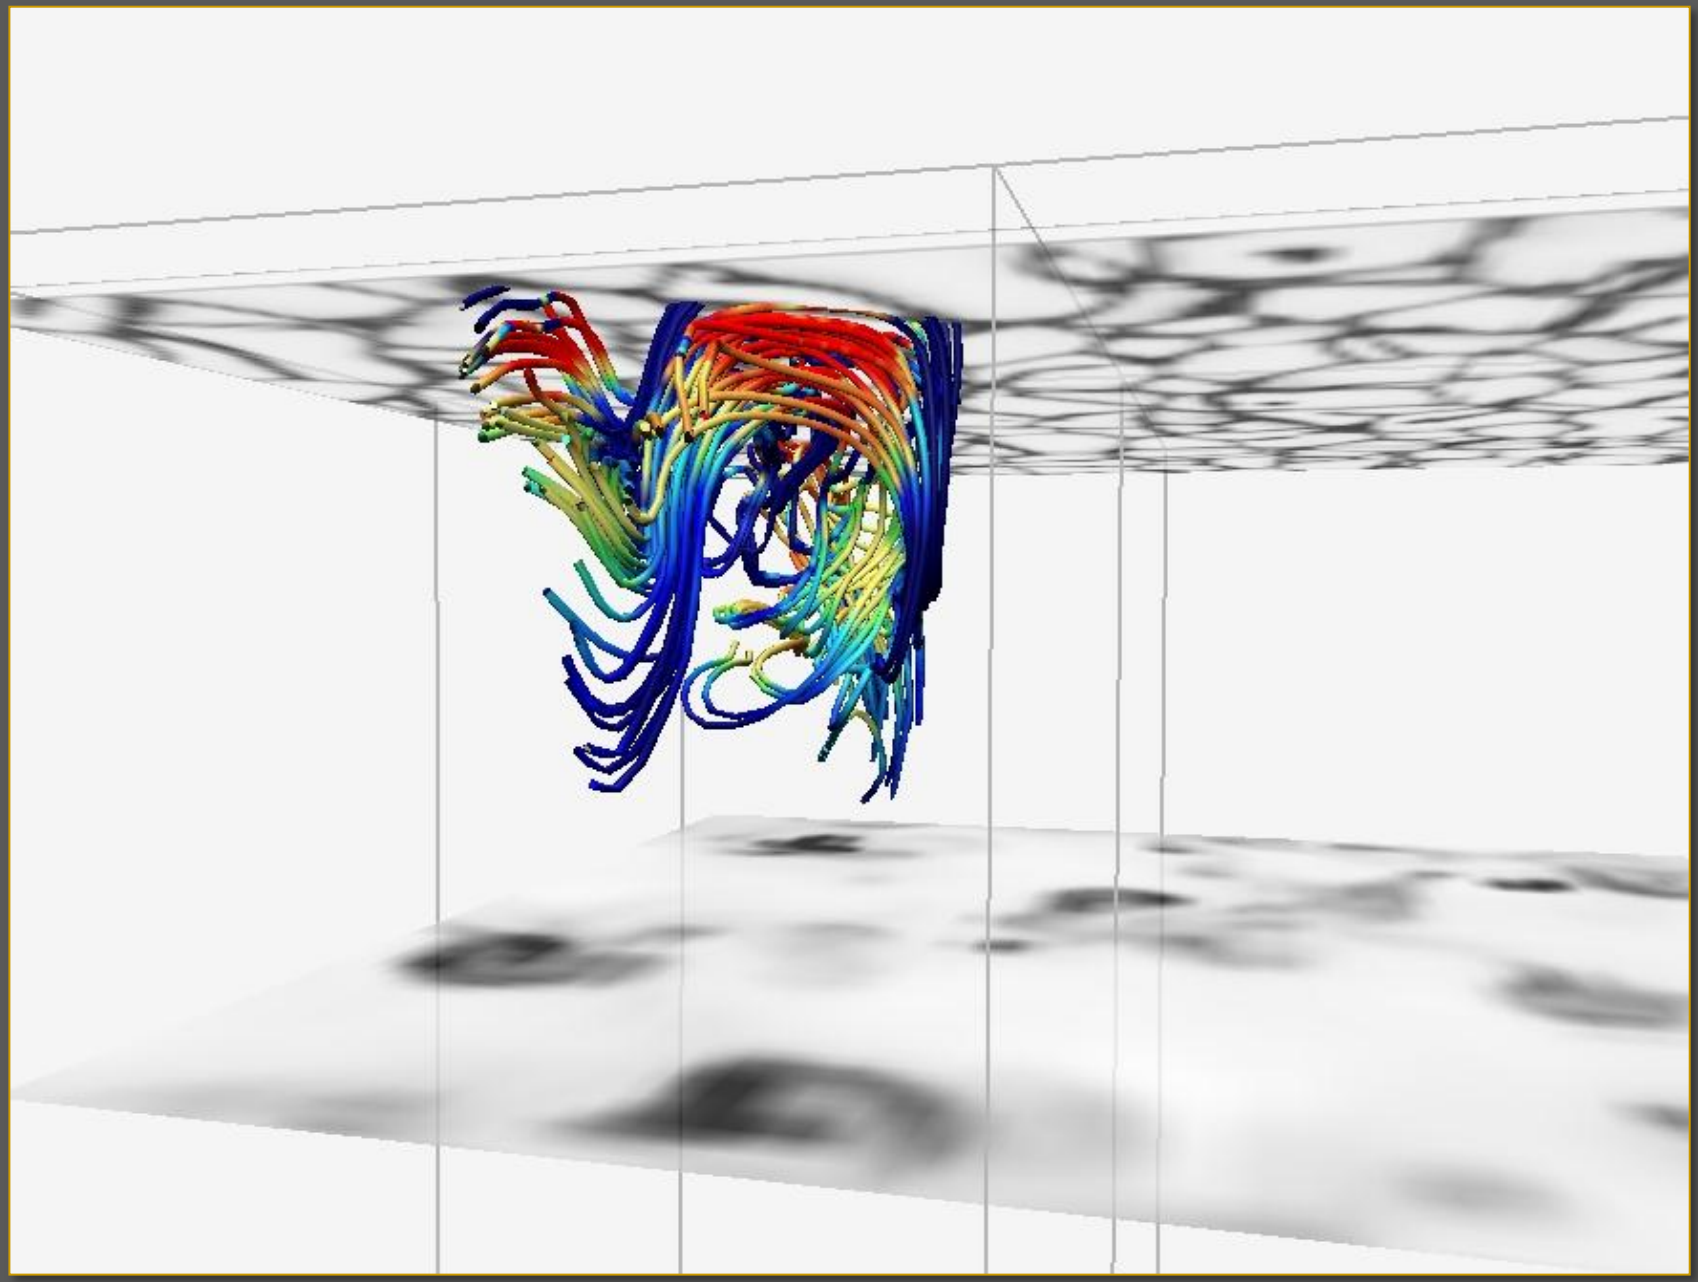

View from above

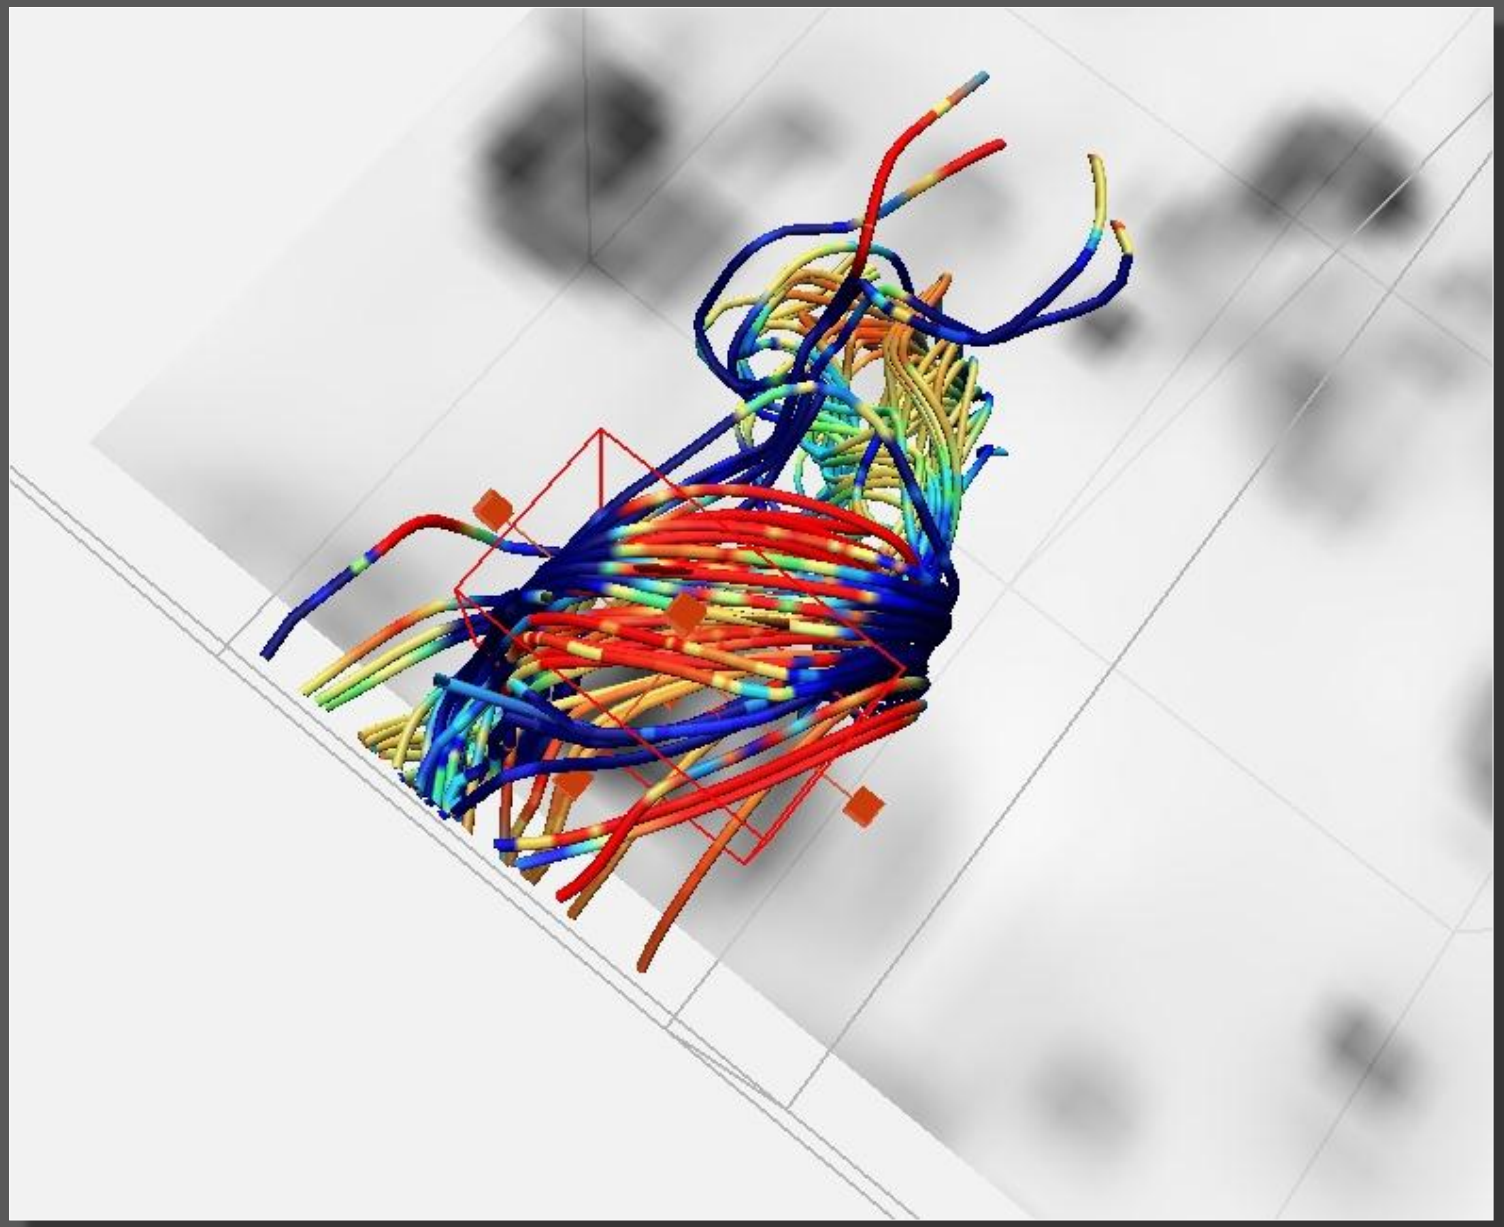

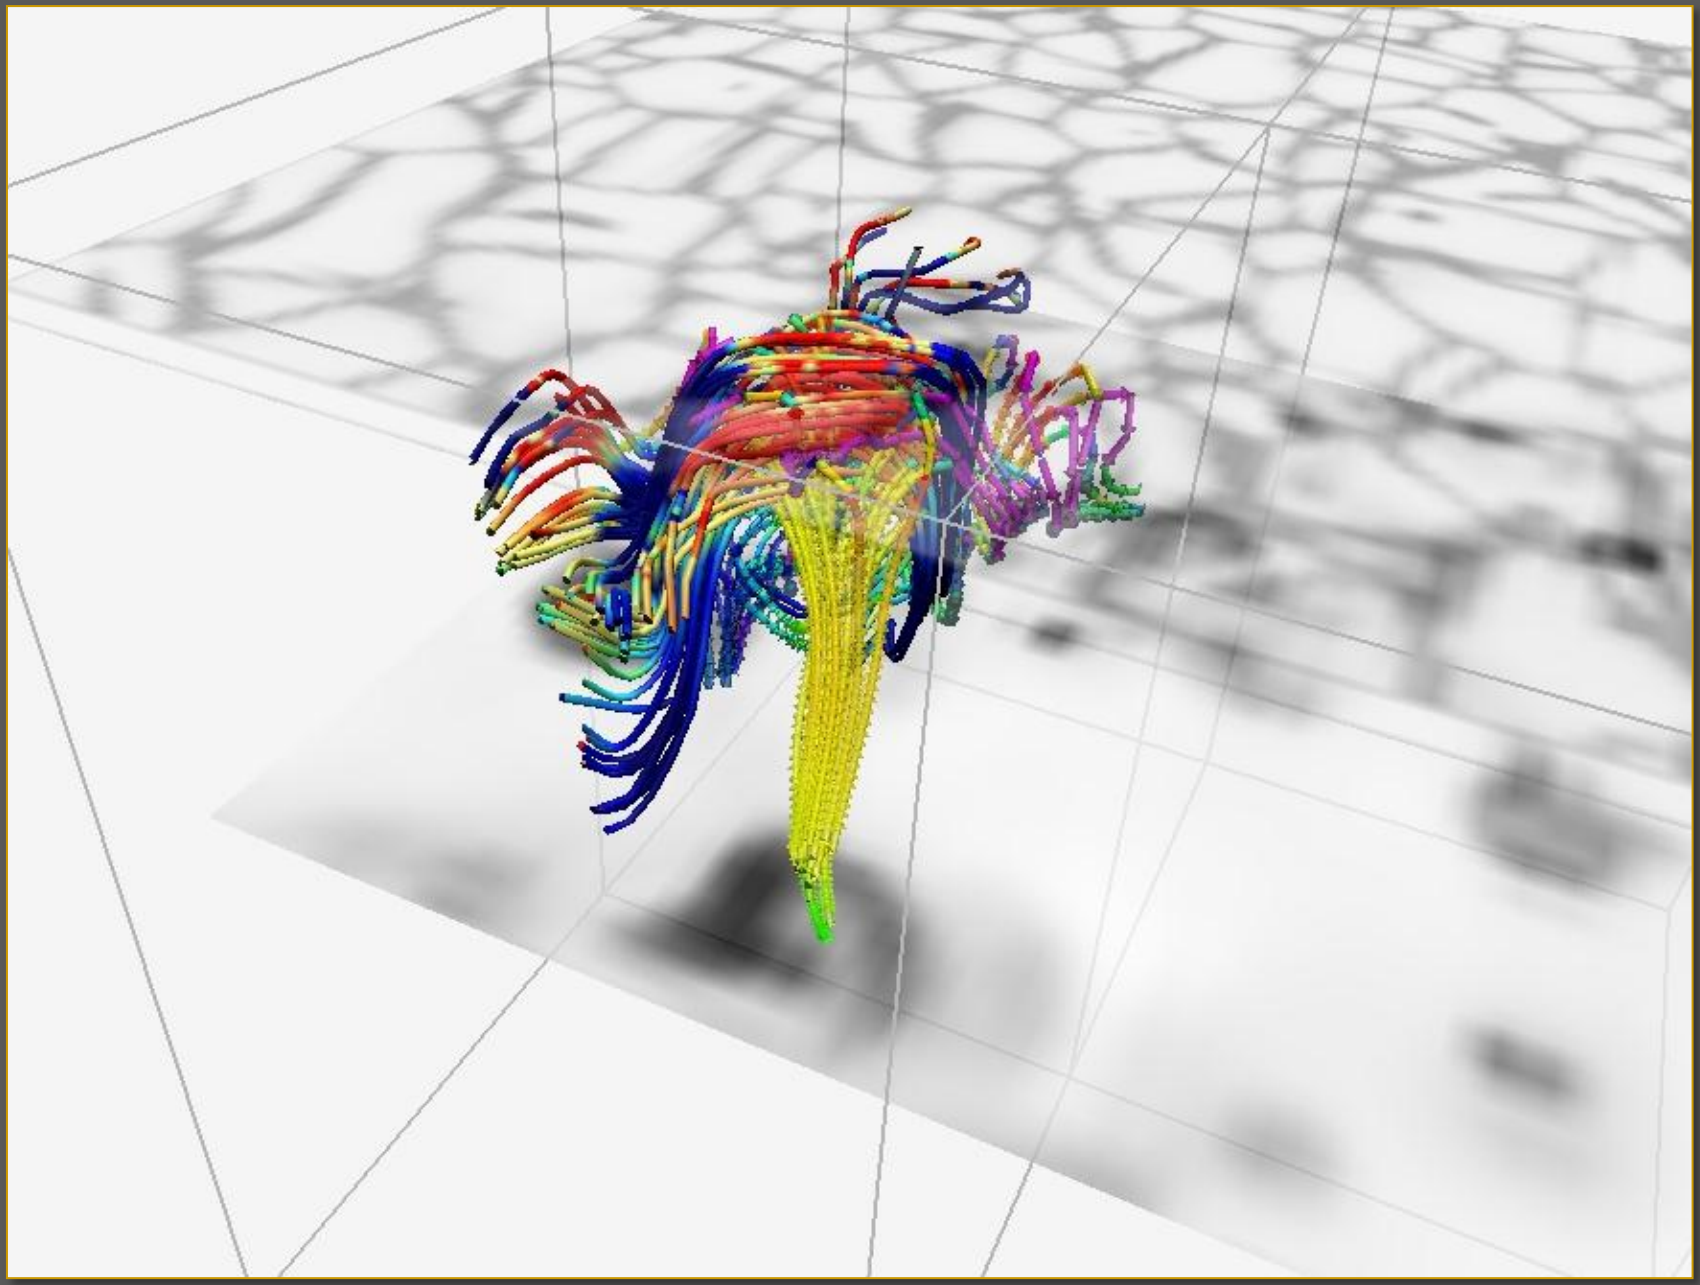

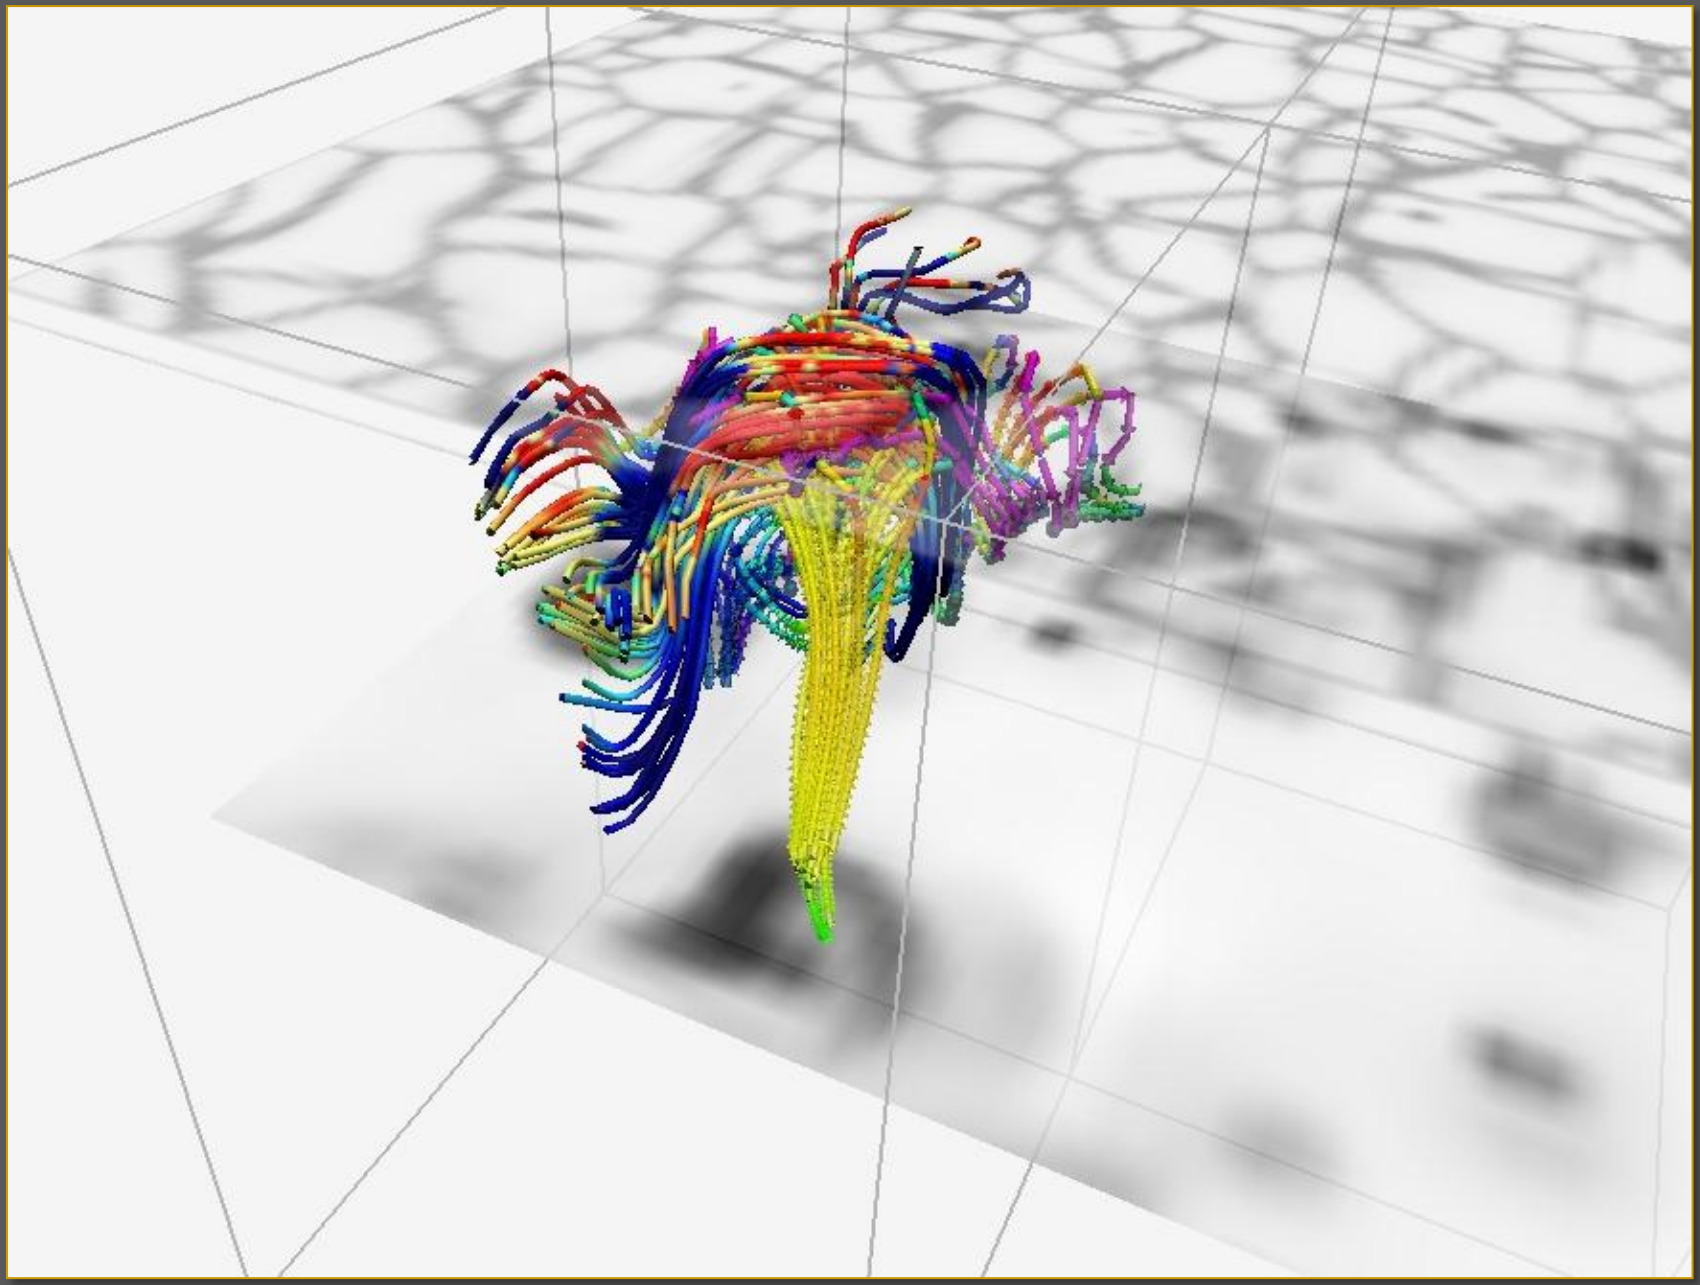

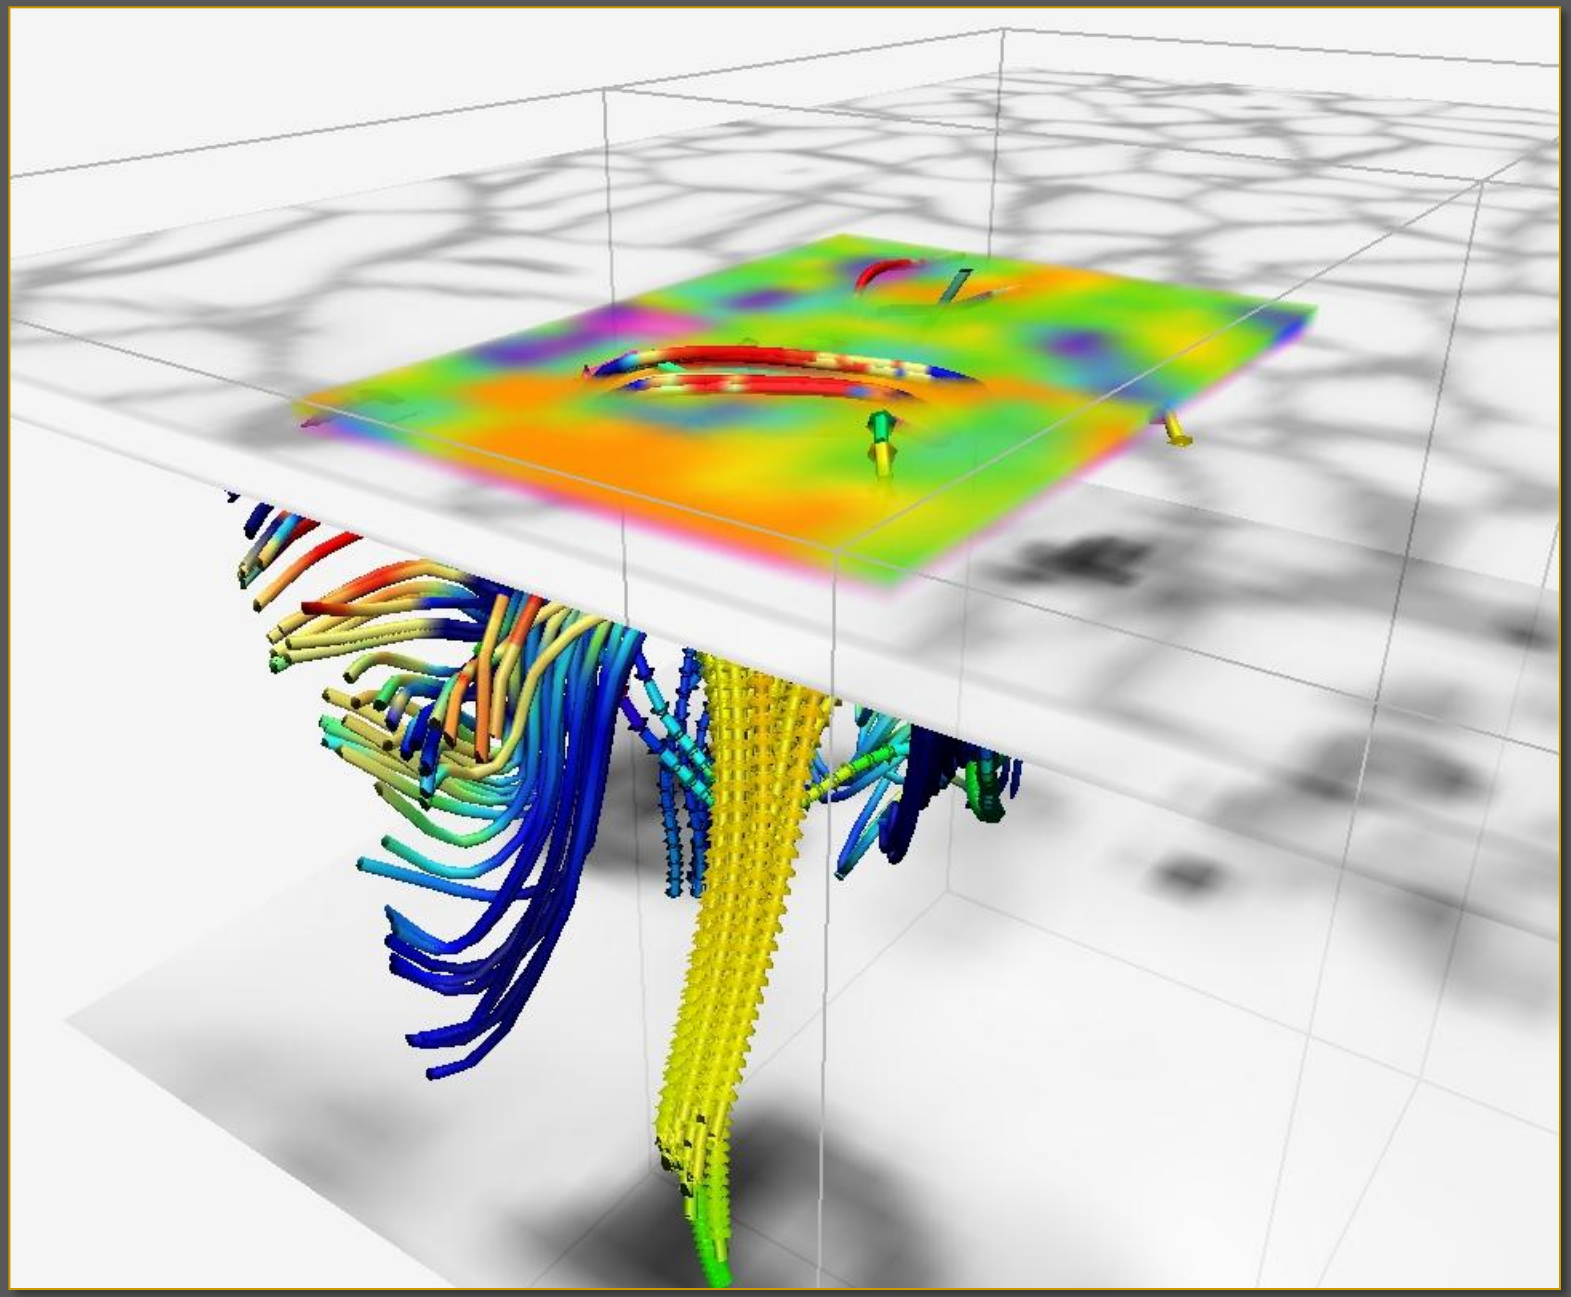

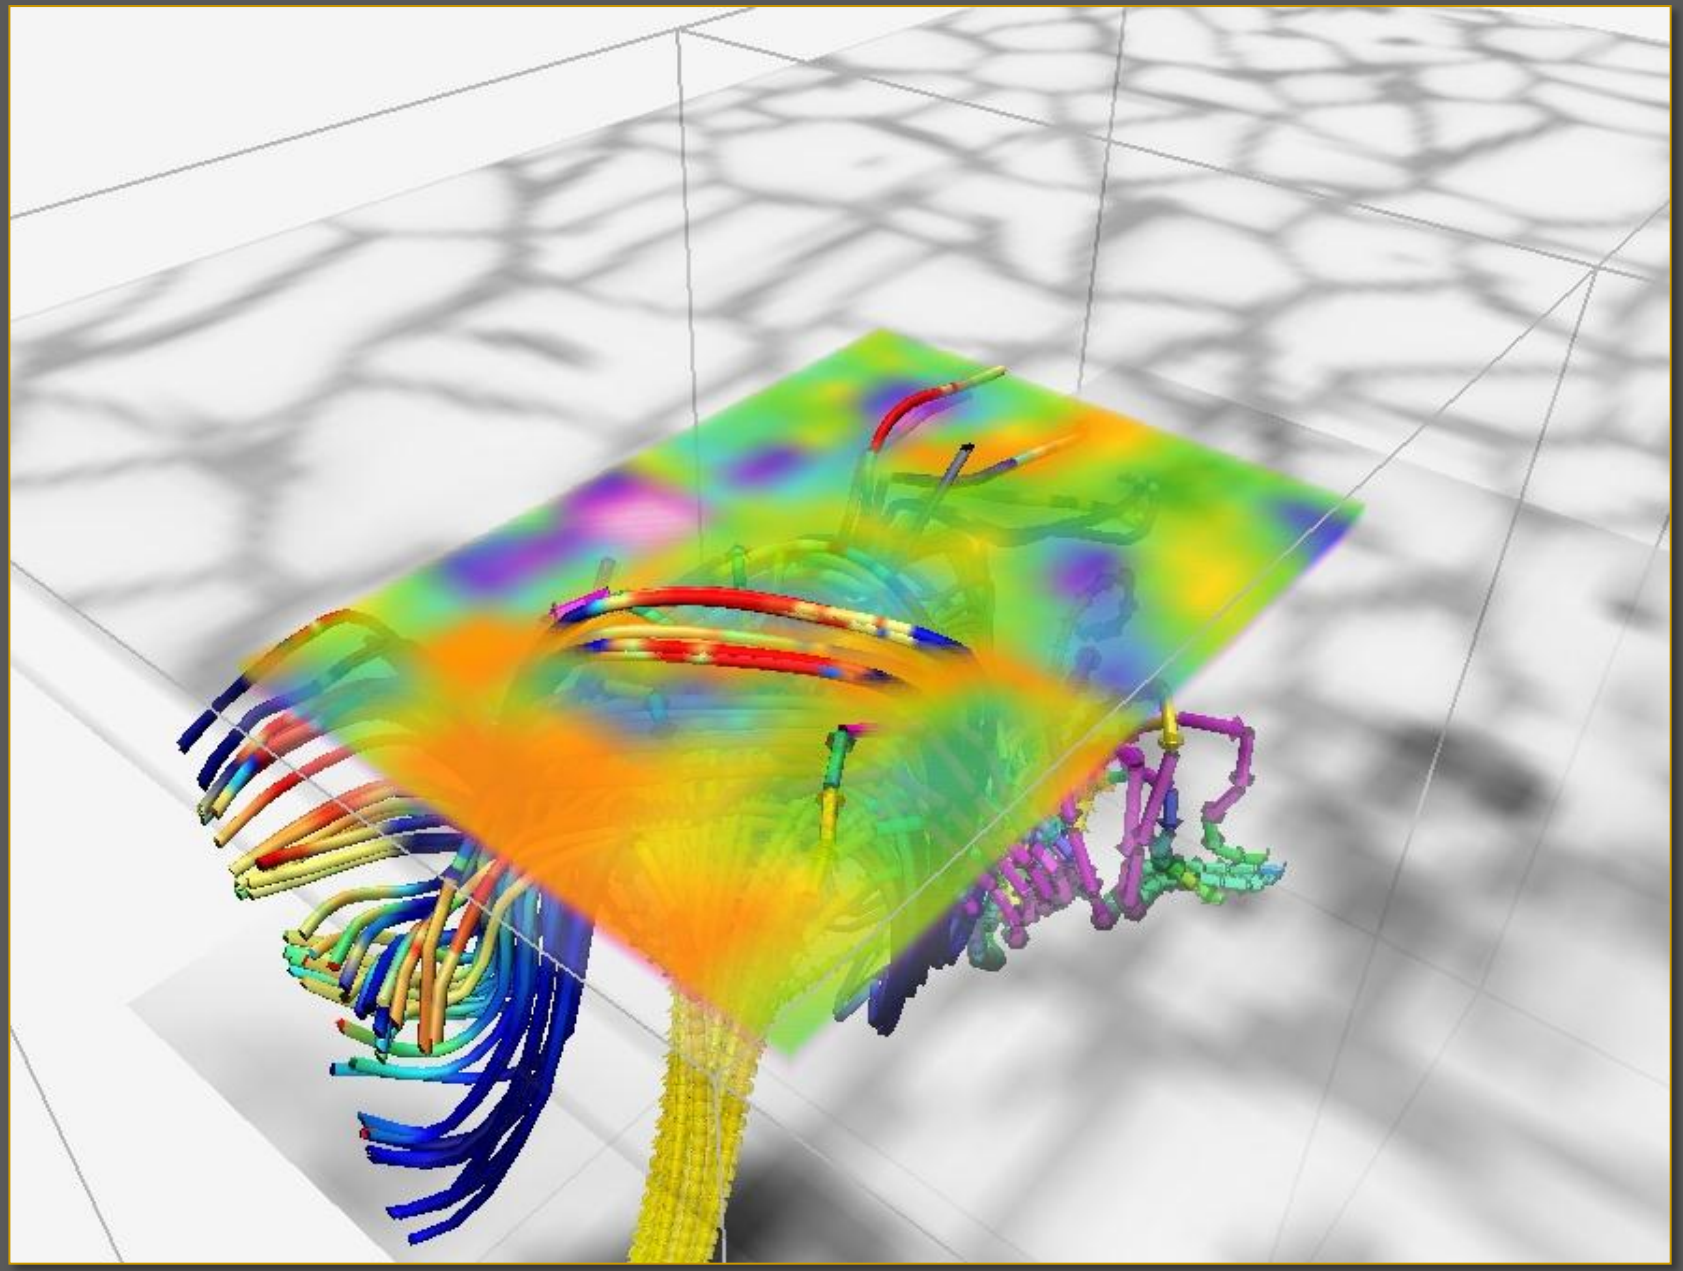

Record Breaking: 135 billion particles

- Solar cut-out, potential field extrapolation, MHD model, particle-model-cutout, particle acceleration

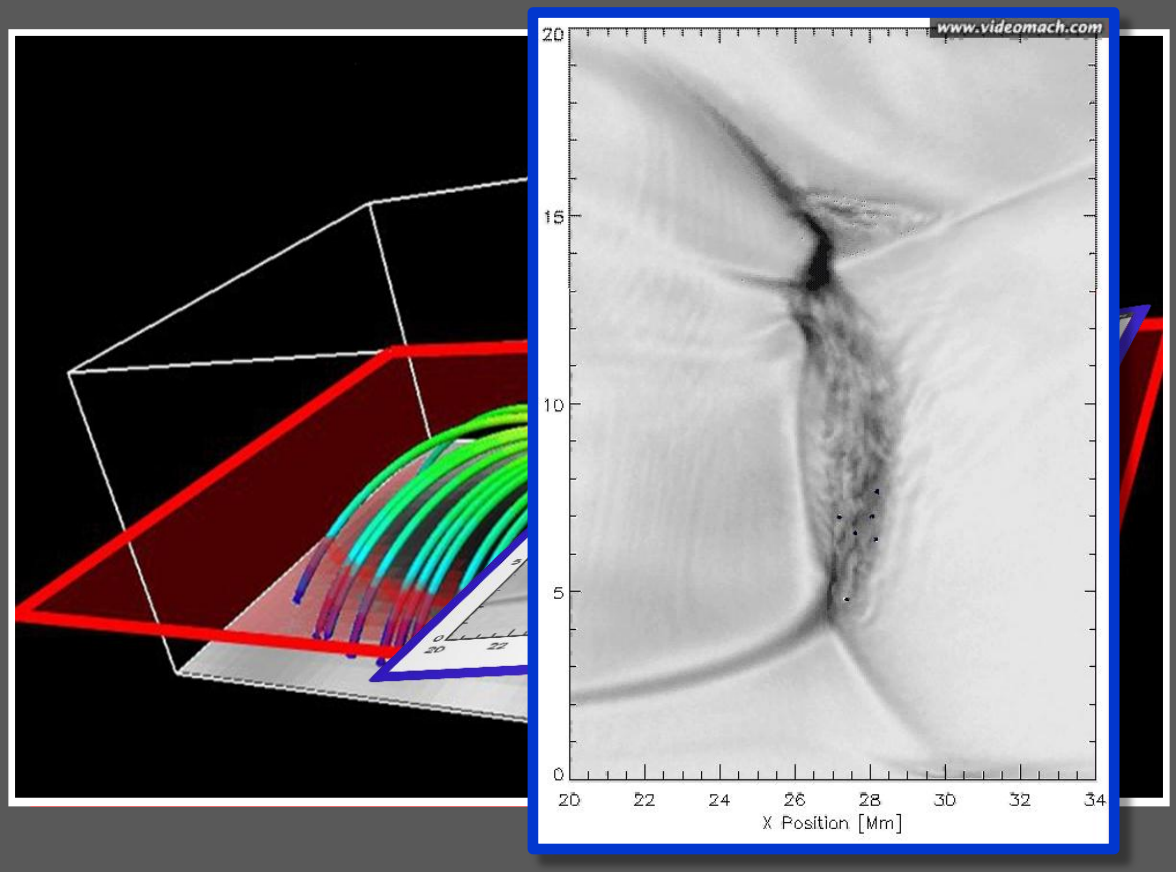

PRACE Specifics

- Intermittent Tier-0 access typical
  - 1 year projects (until now)
    - multi-year (2 yr!) access from 5th call
  - Needs to be complemented by Tier-1 & local facilities
  - Development, training, …
- Front line research generates TB – PB of data
  - Needs to be ***analyzed***, moved & ***archived***
  - ***Access*** should be made possible for peers
    - Sub-projects, derived data …

Analysis Requirements

- Analysis
  - Huge RAM size
  - Multi-core
  - High I/O bandwith

    $\Rightarrow$ Virtual Shared Memory (ScaleMP) is ideal

- At DCSC-KU
  - ScaleMP system, similar to 'Gordon' at San Diego (SDSC)
  - So far: 4 nodes = 48 cores, 256 GB ram, 200 TB disks
  - Expandable ….

Archival Requirement

- PRACE related
  - Data to be moved off PRACE system 2 mo after project end
- Simulation pipe-line – *must* have an open end!
  - Runs (PRACE, NASA/Ames, DCSC-KU)
  - Analysis (local – at FZ-J, Ames, DCSC-KU)
  - Restarts, continuations, spin-off, …
  - Peer access – by collaboration only, so far
  - Archival – Hierarchical Storage Management preferred
    - Absolutely necessary to sustain the pipe line
    - Unfortunately not yet implemented at DCSC-KU

Peer Access

- Collaborations (or not), sub-projects, spin-off, …
  - Should / must allow peers to login, use ***analysis*** facilities
  - This avoids massive waste of bandwith & human time

- Easy to combine with existing analysis facilities!

Summary

- PRACE specifics
  - Intermittent (1-2 yr) Tier-0 access typical
    - Needs to be complemented by Tier-1 & local resources
      - Development, training, …
  - Front-line science ⇒ TB – PB of data generated
    - Needs to be moved & archived

- PRACE-related requirements
  - Local (national?) data facilities
    - ***Analysis*** (virtual shared memory preferred)
    - ***Archival*** (Hierarchical Storage Management preferred)
    - ***Access***: external ***analysis access by peers***