

Scalable Data Schemas for Effective and Efficient Management of Simulated Ground Motion Datasets

A PEER – LBNL workshop

Mar 24, 2025

Houjun Tang
Scientific Data Division
Lawrence Berkeley National Laboratory


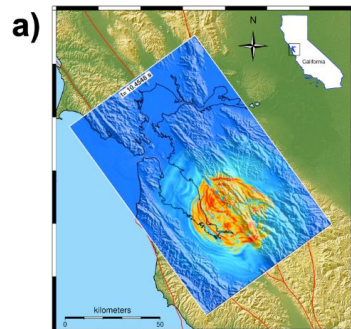



U.S. DEPARTMENT OF
ENERGY

Office of
Cybersecurity, Energy Security,
and Emergency Response




Data management in the EQSIM workflow





Oak Ridge Leadership
Computing Facility (ORNL)



Data Transfer

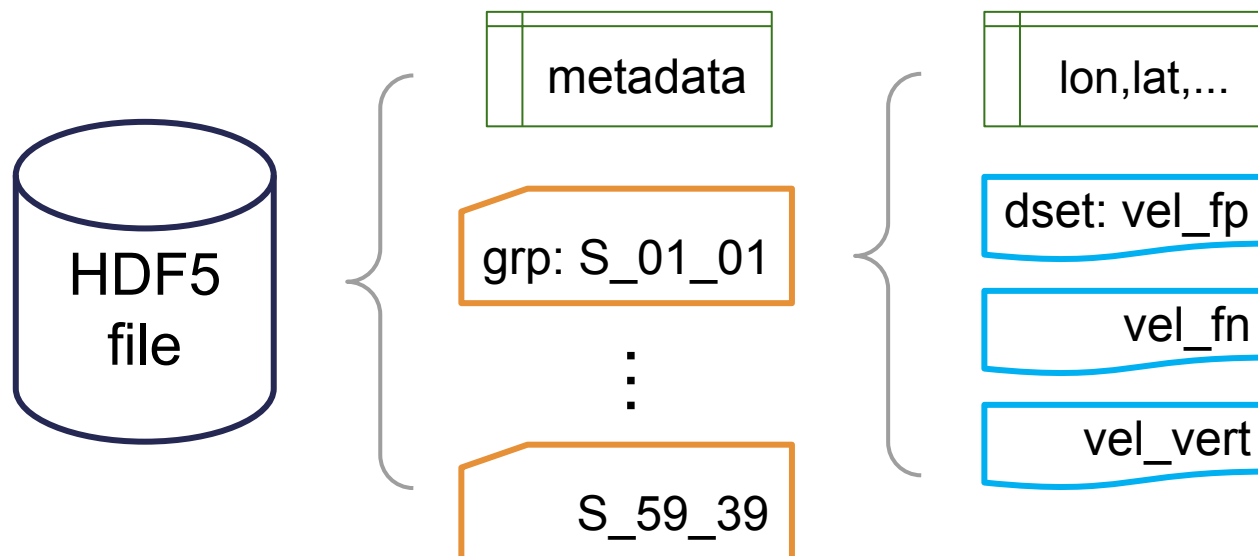


National Energy Research
Scientific Computing Center (LBNL)



Managing EQSIM data with HDF5

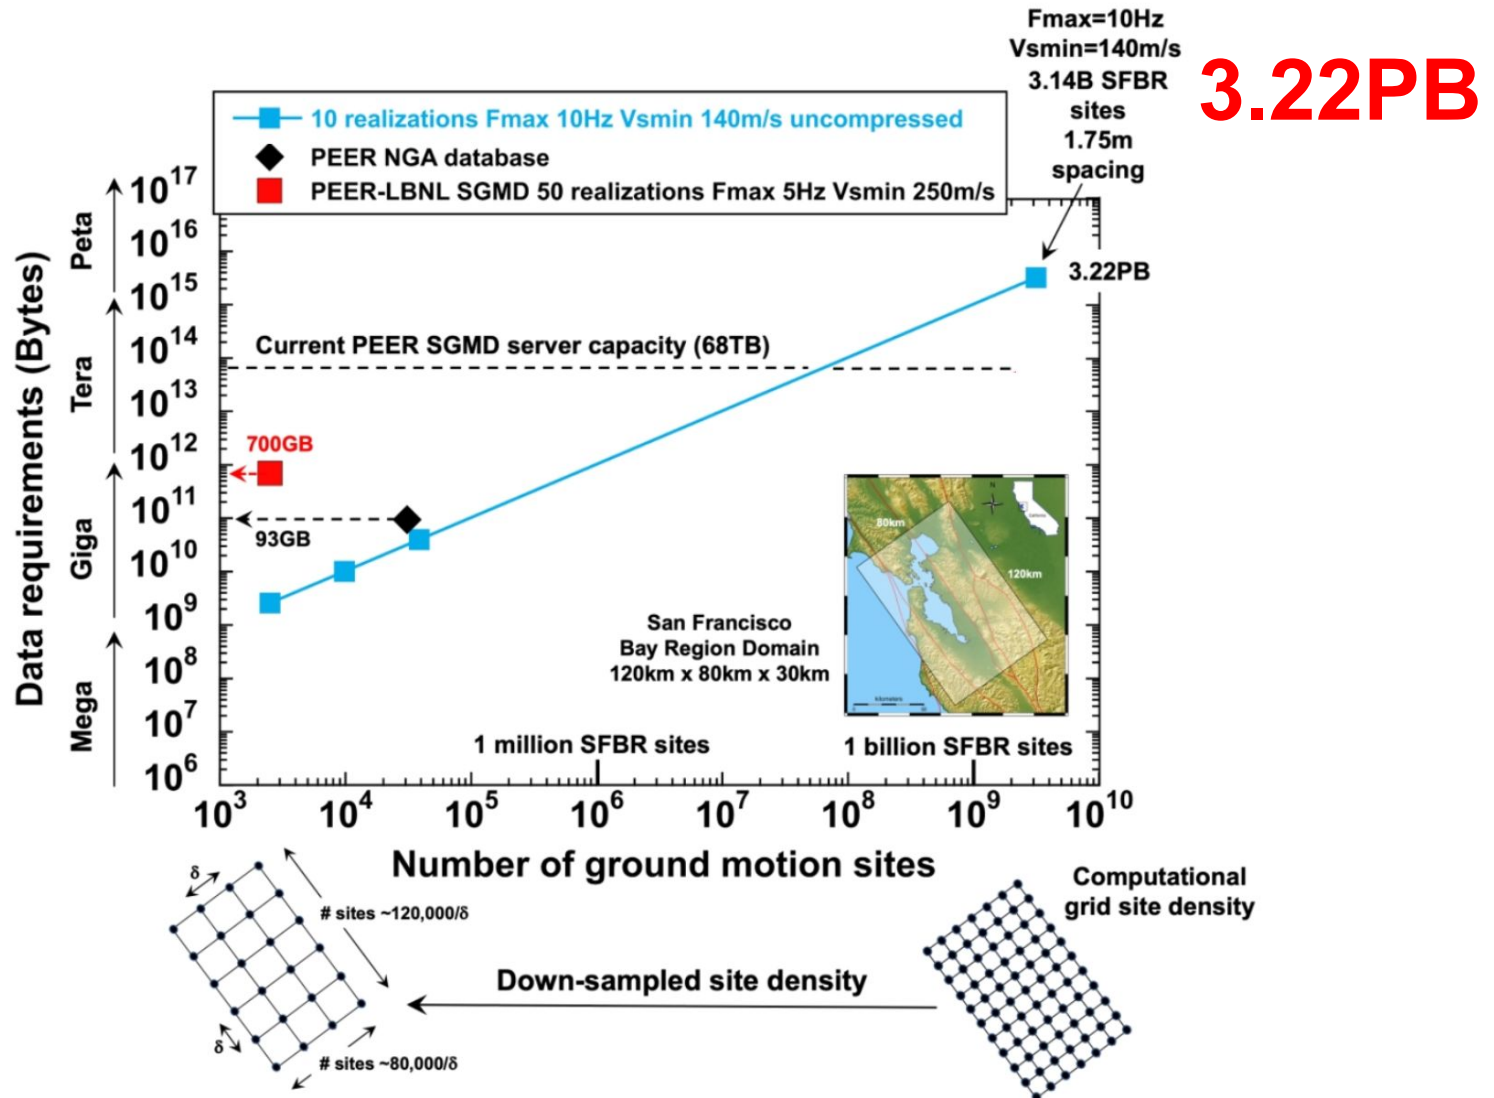
- **HDF5** (Hierarchical Data Format v5) is a data model, library, and file format for managing large and complex scientific data.
 - Fast parallel I/O, easy cross platform sharing, keep metadata with data.
 - Maintained for over 25 years and widely adopted by the scientific community and the industries.



Challenges in EQSIM data management

- **5Hz** Hayward Fault simulations, $V_{s_{\min}}=250$, 120x80km domain
 - 512 Frontier nodes, 6 machine hours for 90s simulation time
 - **2km** spacing, 2,301 x {FP, FN, Vert} ground velocity motions
 - dt=1.2ms, ~**2GB** each, **110GB** total for 50 runs
 - SGMD >**700GB** total of {vel,acc,disp}x{FP,FN,Vert} in ASCII format
 - **6.25m** spacing, **246 million** x {FP, FN, Vert} ground velocity motions
 - Downsample 16x, dt=19ms, ~**25TB** each, **1.2PB** total for 50 runs
- **10Hz** Hayward Fault, $V_{s_{\min}}=250$
 - 3072 Frontier nodes, 10 machine hours for 90s simulation time
 - **3.125m** spacing, **1 billion** x {FP, FN, Vert} ground velocity motions
 - Downsample 15x, dt=9ms, **208TB** each

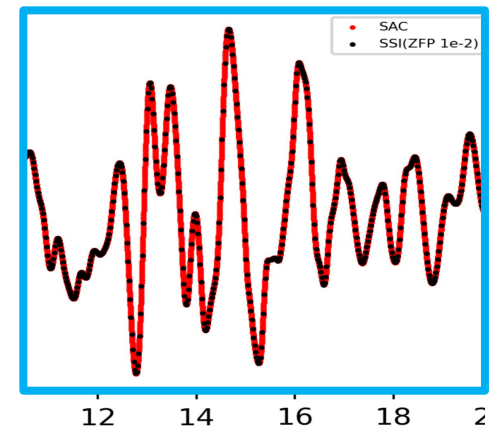

Data becomes impractical to store



We need data reduction

■ Lossless vs lossy compression

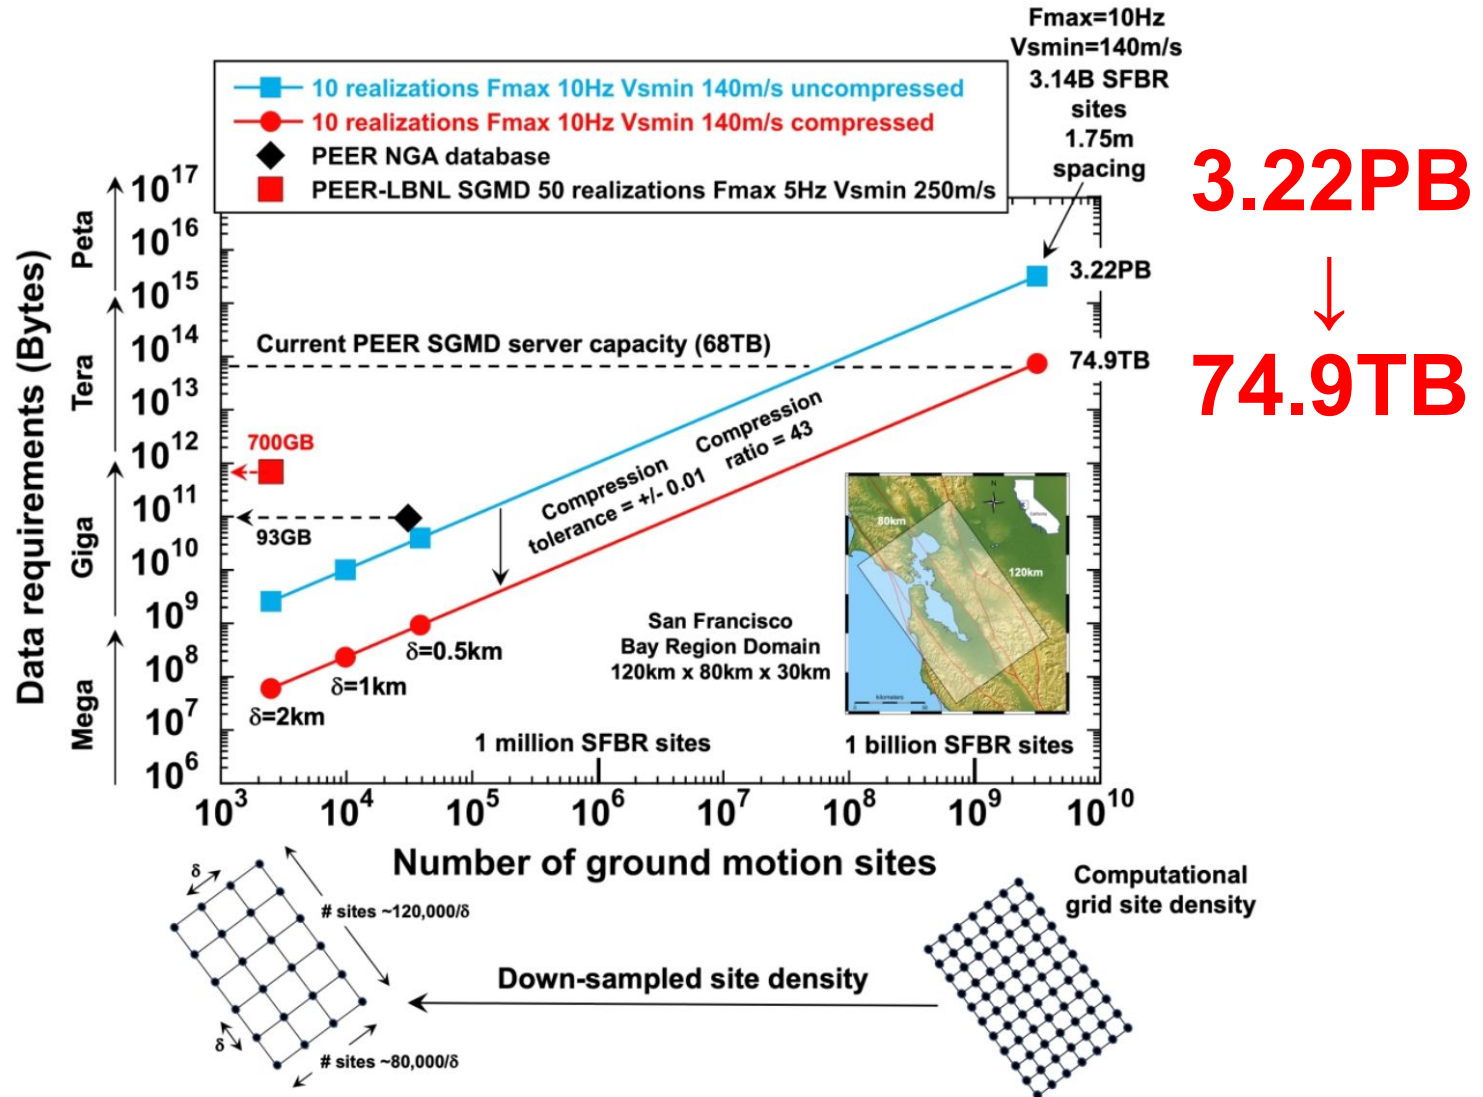
- Split data into small “chunks”, often exploits similarity in values of nearby elements
 - [1.401, 1.419, 1.431, 1.445]
 - find deltas from 1.401 [0, 0.018, 0.03, 0.044]
 - store the deltas with compact encoding to reduce the size, 32→20 bytes
- Lossless preserves all information, $1.401 \rightarrow z_0 \rightarrow \mathbf{1.401}$
- (Error-bounded) Lossy discards some information, $1.401 \rightarrow z_{0.01} \rightarrow \mathbf{1.407/1.399/...}$
- Compression ratio: lossless often < 2 , lossy can be **10 to 100**



Data reduction through lossy compression

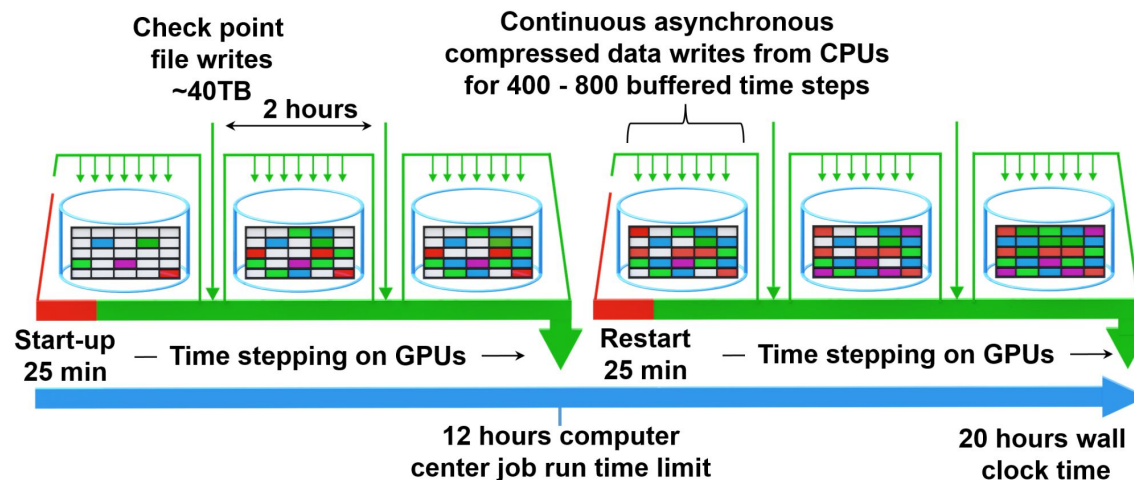

- **5Hz** Hayward Fault simulations, $V_{s_{\min}}=250$, 120x80km domain
 - **6.25m** spacing, **246 million** x {FP, FN, Vert} ground velocity motions
 - ZFP-compressed HDF5 file, error bound=0.01, dt=19ms
 - ~**25TB** compressed to ~**290GB**.
 - 50 runs: **1.2PB** compressed to **14TB**.
- **10Hz** Hayward Fault, $V_{s_{\min}}=250$
 - **3.125m** spacing, **1 billion** x {FP, FN, Vert} ground velocity motions
 - ZFP-compressed HDF5 format, error bound=0.01, dt=9ms
 - 1 run: **208TB** compressed to **4TB**

Data reduction through lossy compression

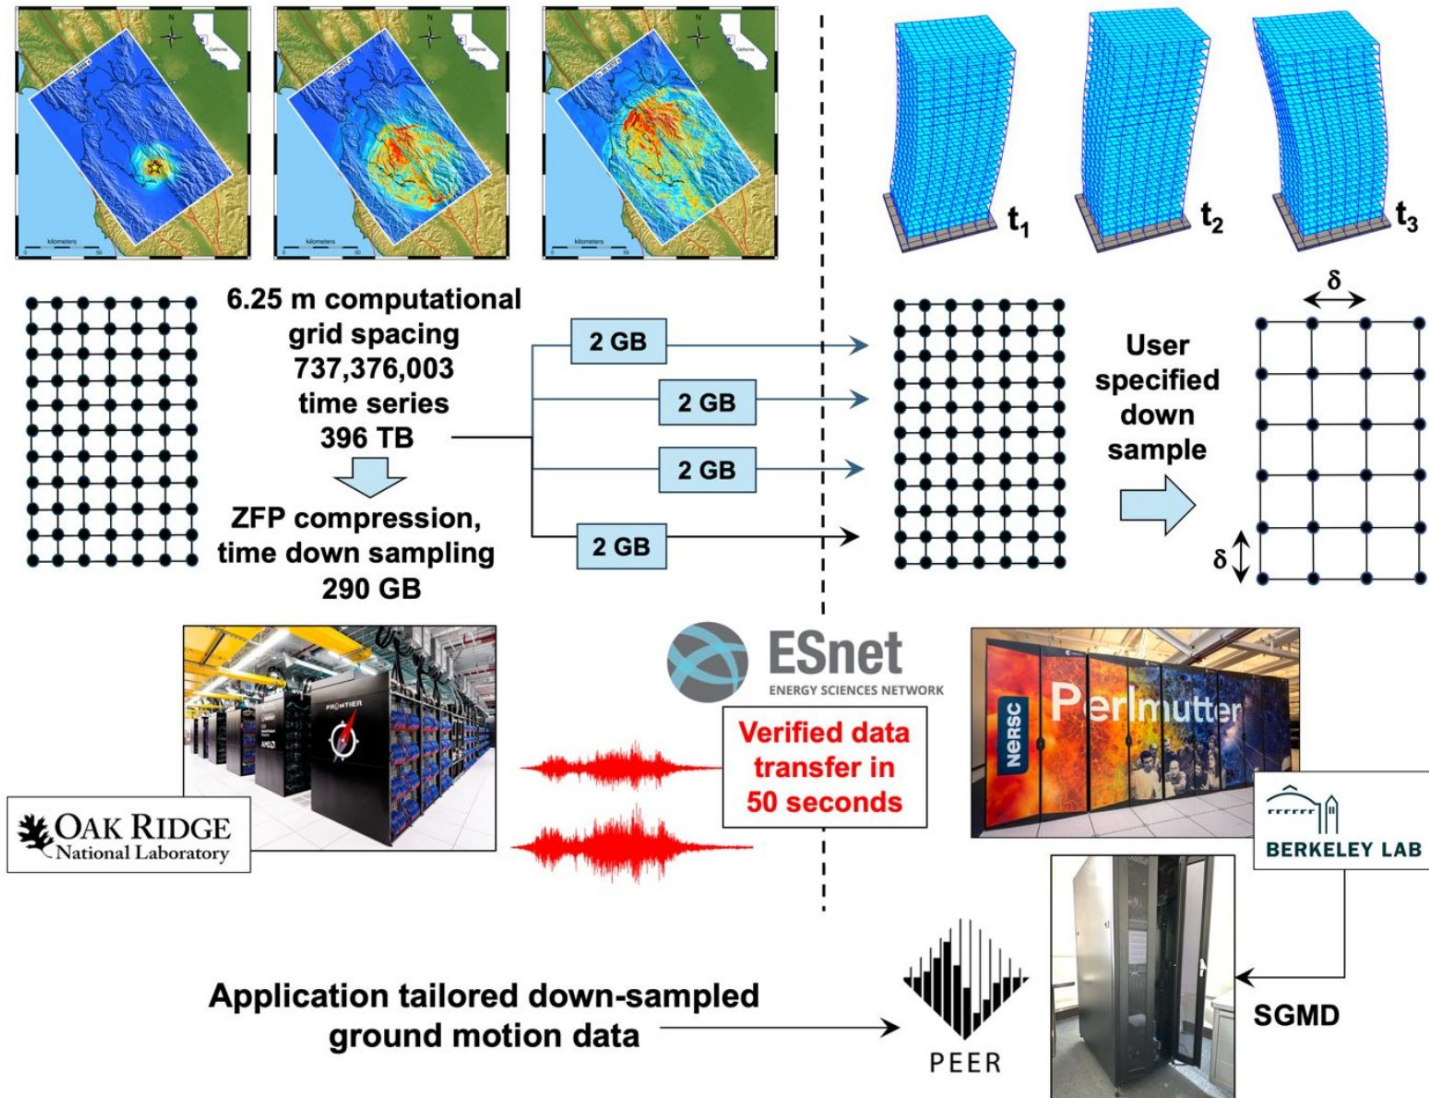


Optimizing parallel data write during simulation

- HDF5 data chunking
 - Required by compression
 - More contiguous writes
- Data buffering
 - Reduce the number of writes
 - More efficient parallel I/O
- Asynchronous I/O
 - Overlaps computation with I/O utilizing idle CPU cores when GPU is computing




Managing data across facilities



Optimizing data read for analysis and visualization


- Data uses write-optimized chunked (t,x,y,z) layout
- Visualization needs individual (x,y) cross-sections
- Serial data access is too slow
 - Original serial code to generate animation takes several **hours**
- Efficient parallel data access must be **chunk-aware**
 - Distribute the workload with chunk-aligned domain decomposition
 - Read entire chunk (instead of row) and cache in memory for data extraction

} **mismatch**




Chunk-aware parallel processing

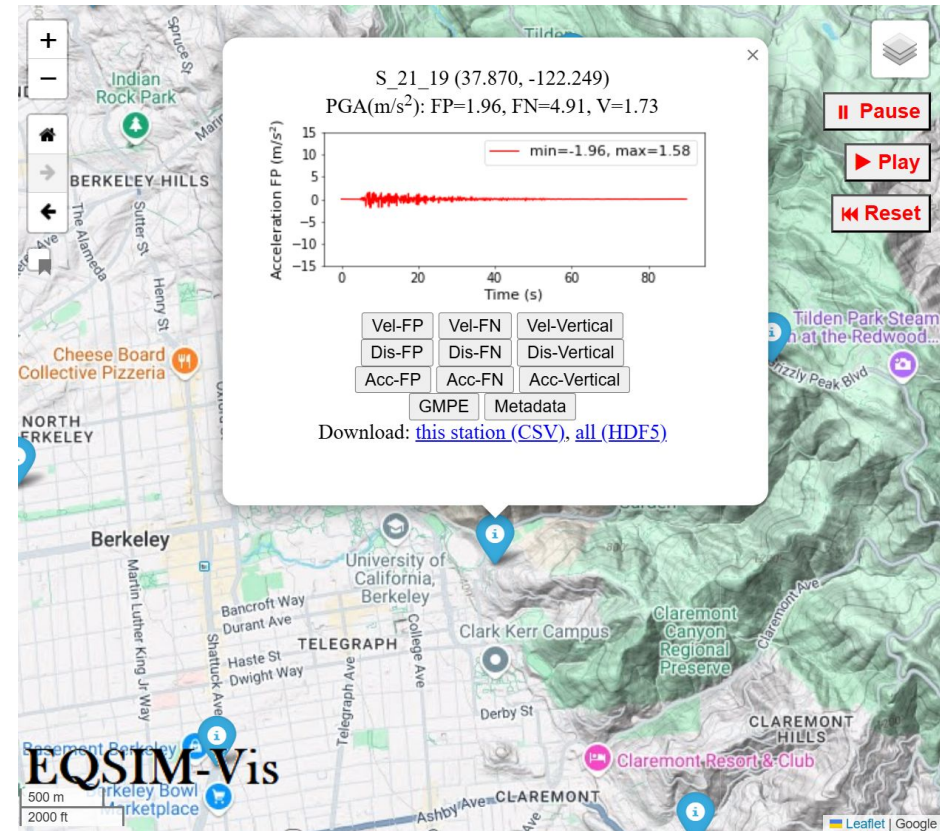
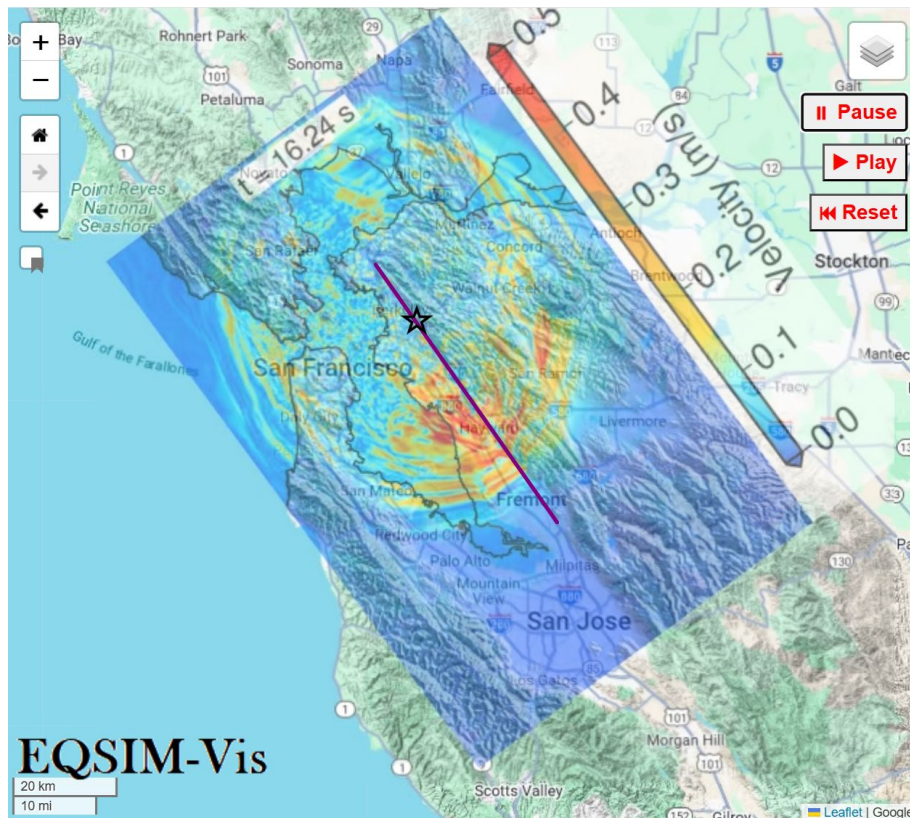
- **2 horizontal components**
- **190 GB** compressed data (ZFP lossy, accuracy=0.01)
- **9 TB** uncompressed
- Downsample in time (2x) space (10x)
- Serial processing takes **10+ hours**



7 minutes on 8 Perlmutter nodes
with chunk-aware processing



EQSIM-Vis: inspect ground motions on an interactive map



Summary

- HDF5 is effective for managing large data
 - Cross-platform, multi-language support (C, Python, MATLAB)
 - Self-describing, stores metadata together with data
 - Efficiently parallel I/O
 - Transparent compression support
- Effective error-bounded lossy compression can significantly reduce the total data size.
- Storage-aware parallel processing is key to efficient data analysis and visualization
- Visualization tools such as EQSIM-vis enables efficient data inspection.

Thanks!

Questions?

email: htang4@lbl.gov