Towards Exascale Simulations for Regional-Scale Earthquake Hazard and Risk

David McCallen
Lawrence Berkeley National Laboratory

BERKELEY LAB

UNIVERSITY OF CALIFORNIA

HPC Users Forum, Tucson Arizona, April 2018

ECP
EXASCALE COMPUTING PROJECT

The Department of Energy Labs have driven the U.S. advancements in scientific HPC

1,000,000,000,000,000,000 (10^{18}) Flops (ECP)

100,000,000 (10^{8}) Flops
(My arrival as a grad student at LLNL)

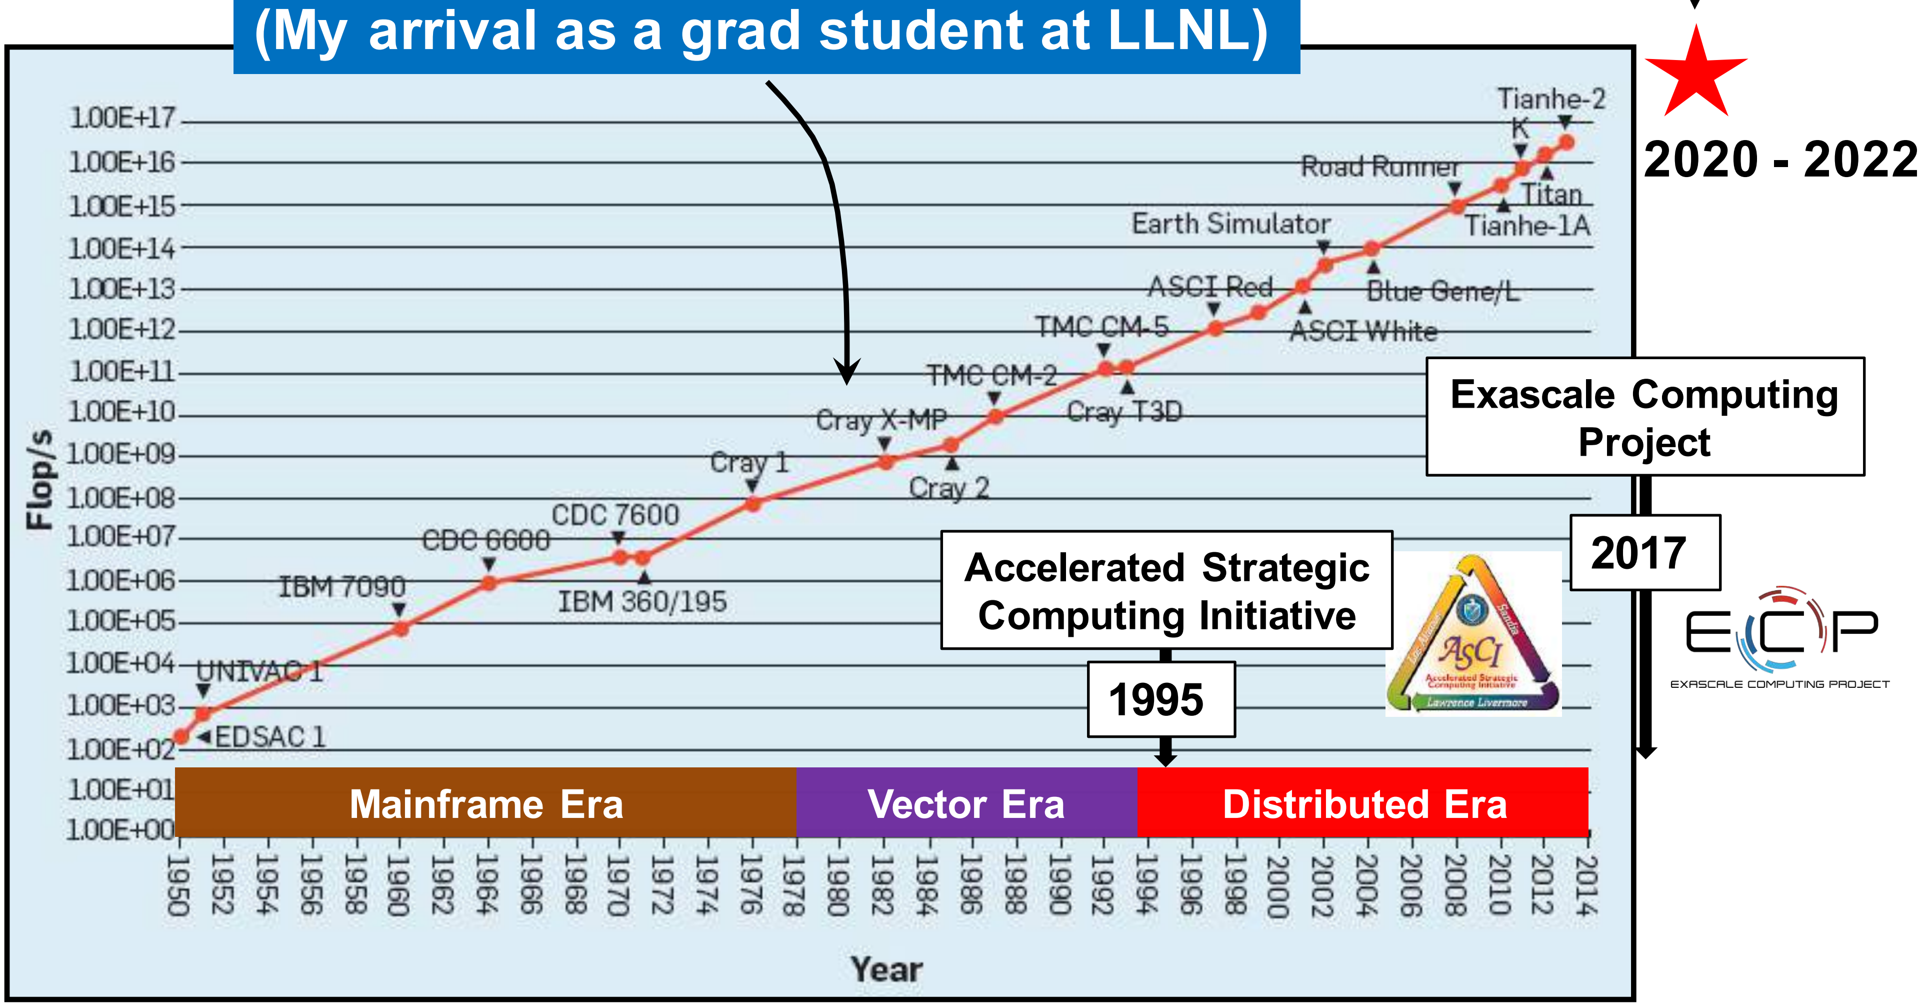

The DOE Exascale Computing Project is a concerted effort to accelerate U.S. HPC

Components:
Advanced hardware development
Application development
Software technology development

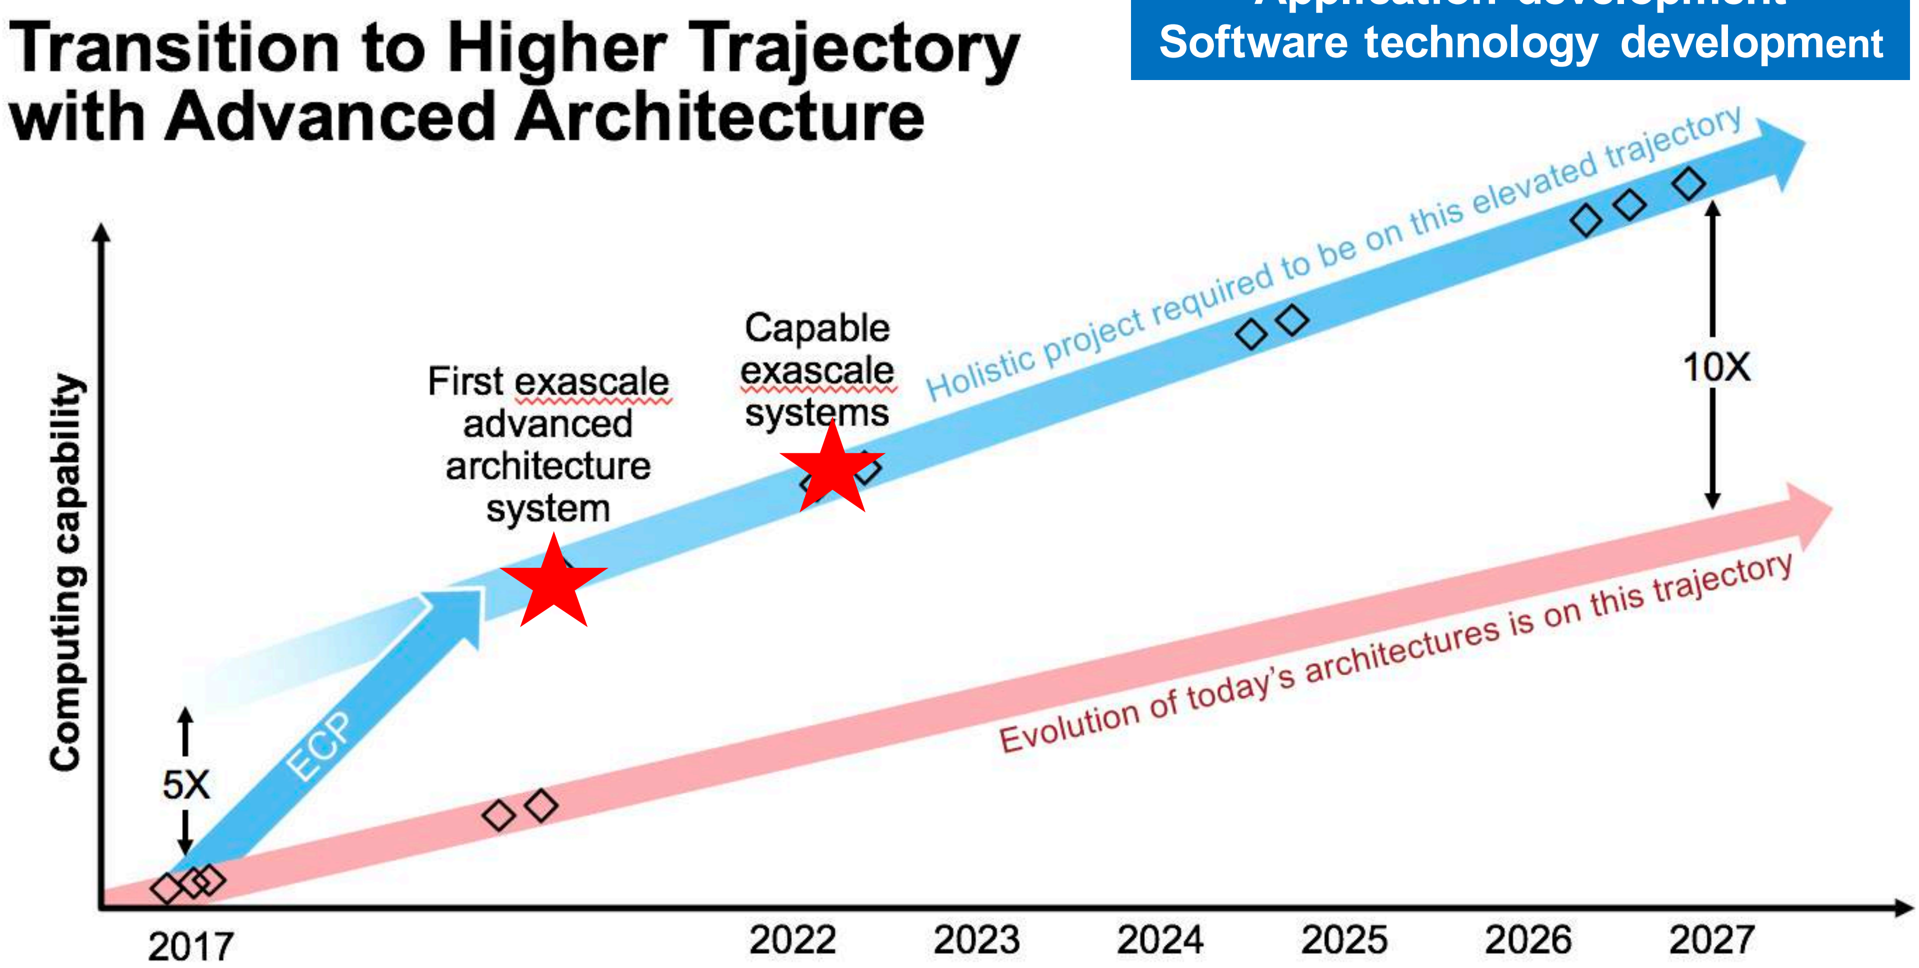

Context - this is a transformational era for earthquake hazard and risk assessments

Historical (1960s - now)
Empirically based

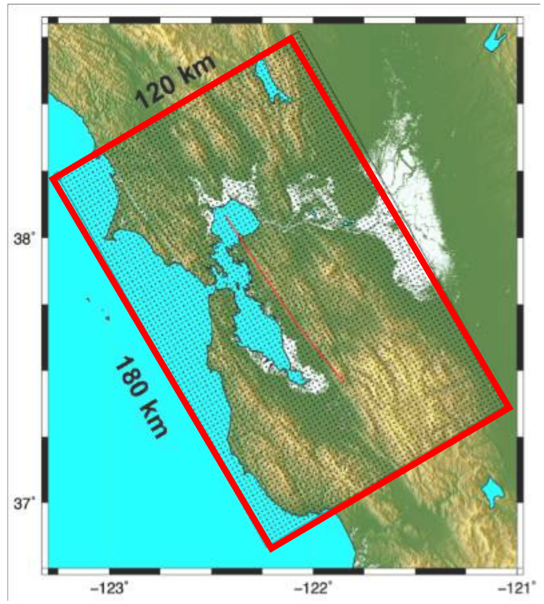

Future (now - forward)
Simulation (physics) based

Probabilistic Seismic Hazard

Simulation Based Hazard & Risk

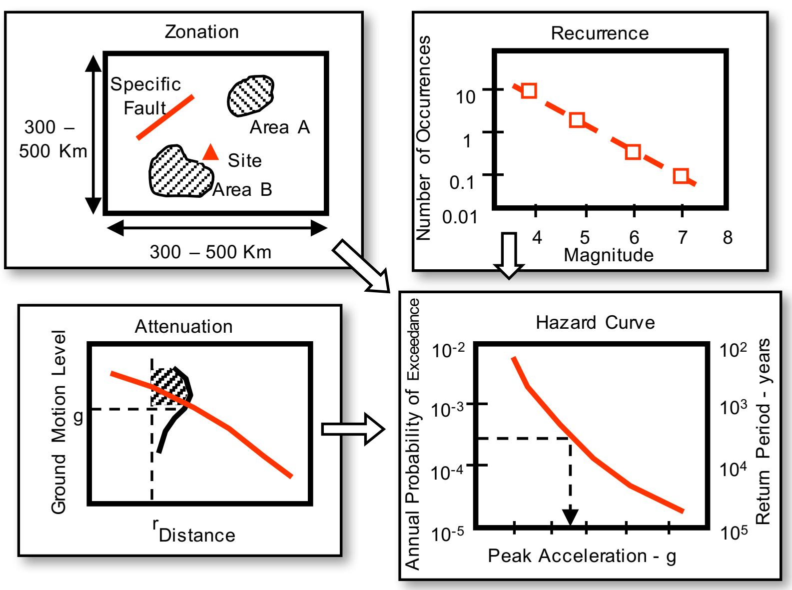

Realistic simulations can both increase understanding and reduce uncertainties

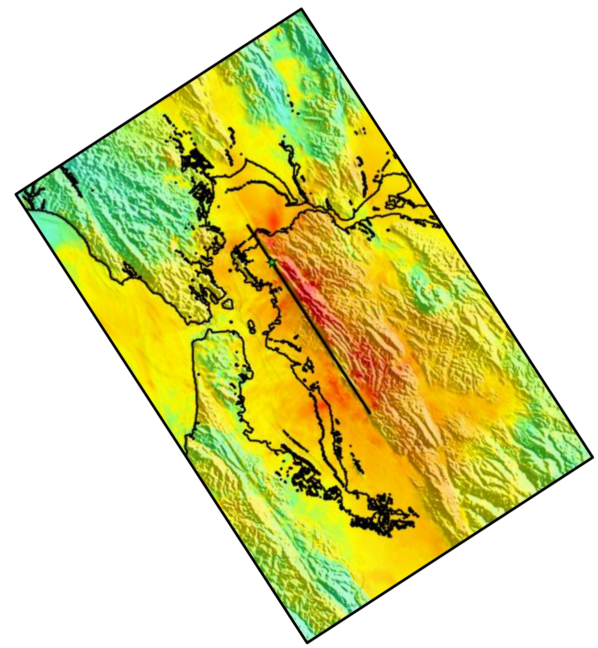

Our ECP objective is a computational framework for earthquake hazard and risk

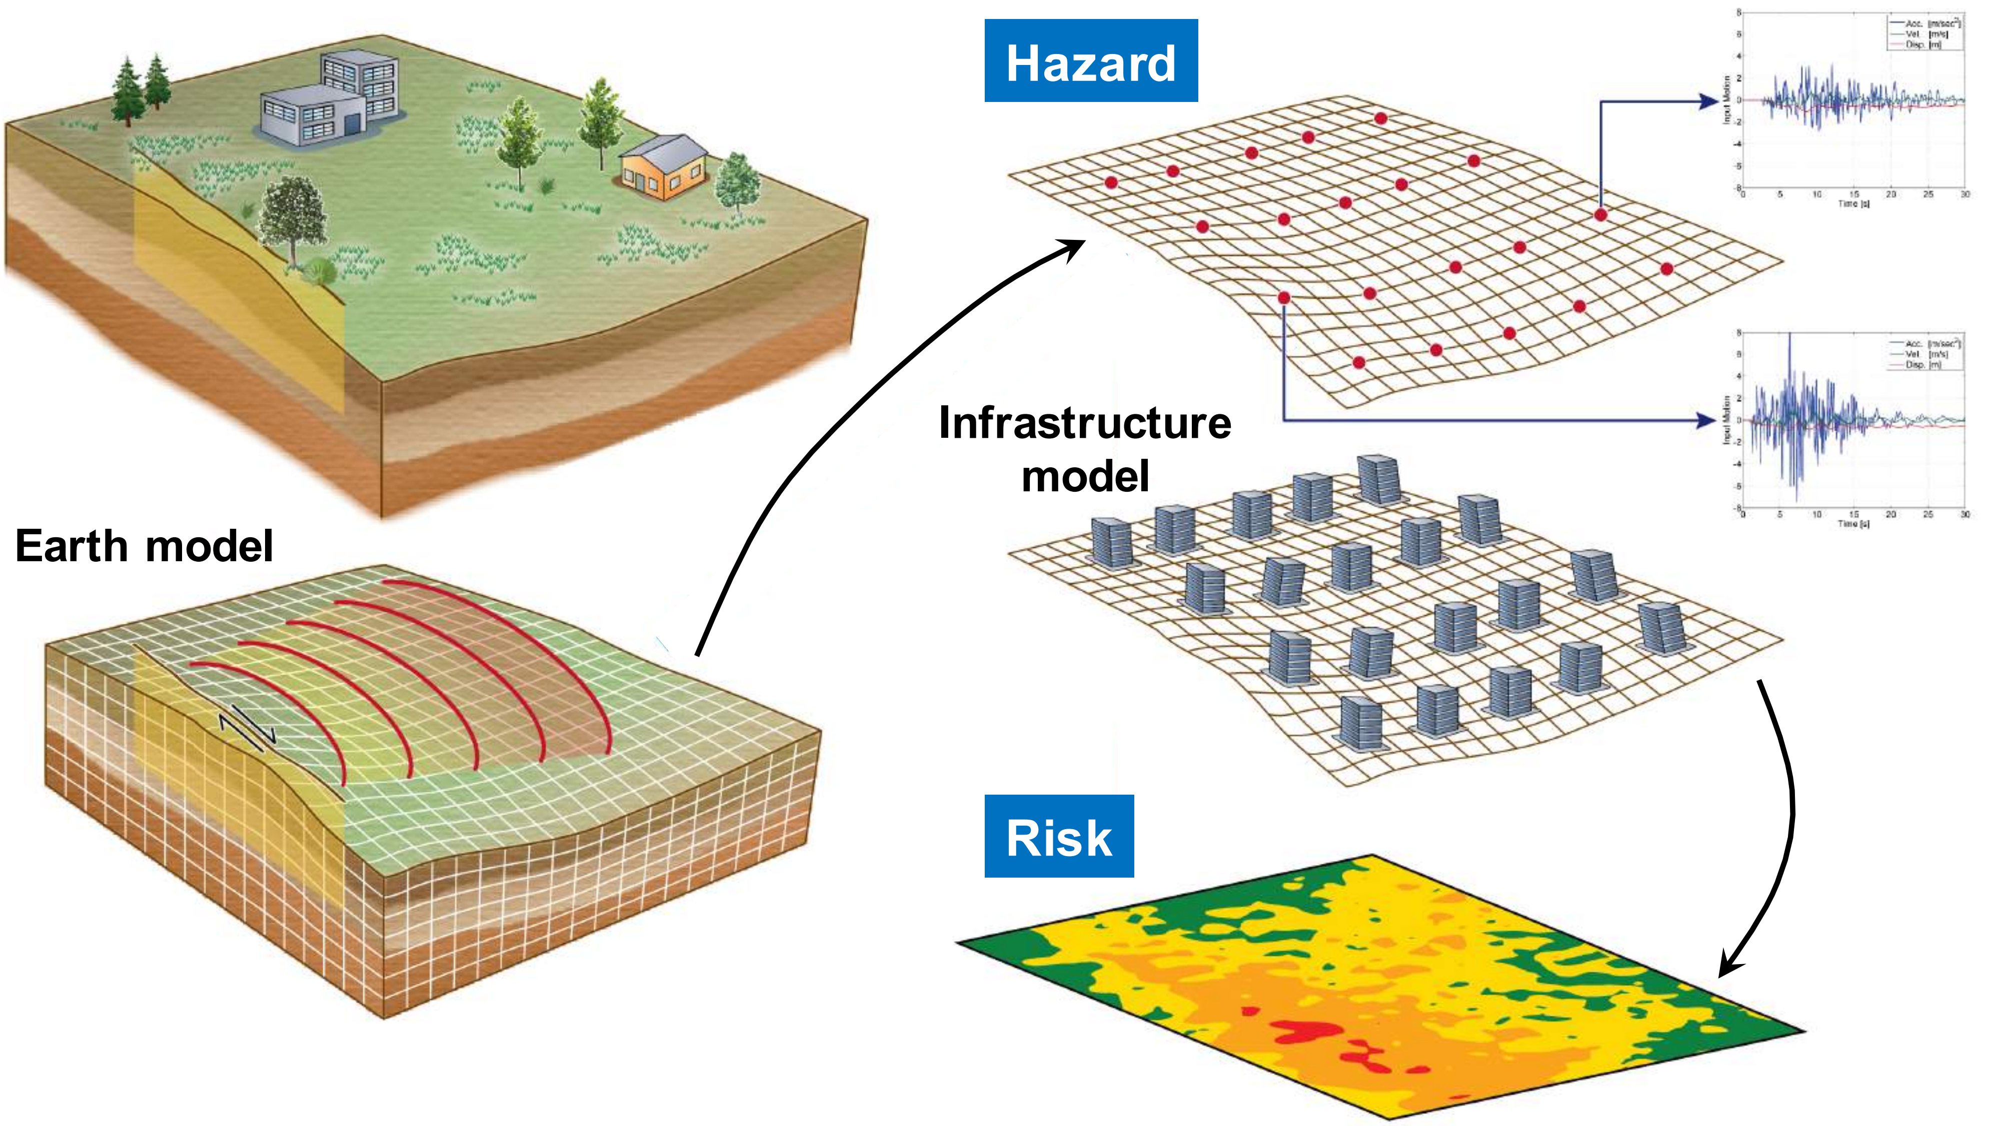

A multidisciplinary team is essential – a National Laboratory scale problem

Computational Science and Applied Math

Dr. Anders Petersson

Dr. Hans Johansen

Structural Mechanics

Dr. David McCallen

Dr. Mamun Miah

Seismology

Dr. Arthur Rodgers

Dr. Arben Pitarka

Geo Mechanics

Dr. Boris Jeremic

We are advancing and coupling codes for geophysics and infrastructure modeling

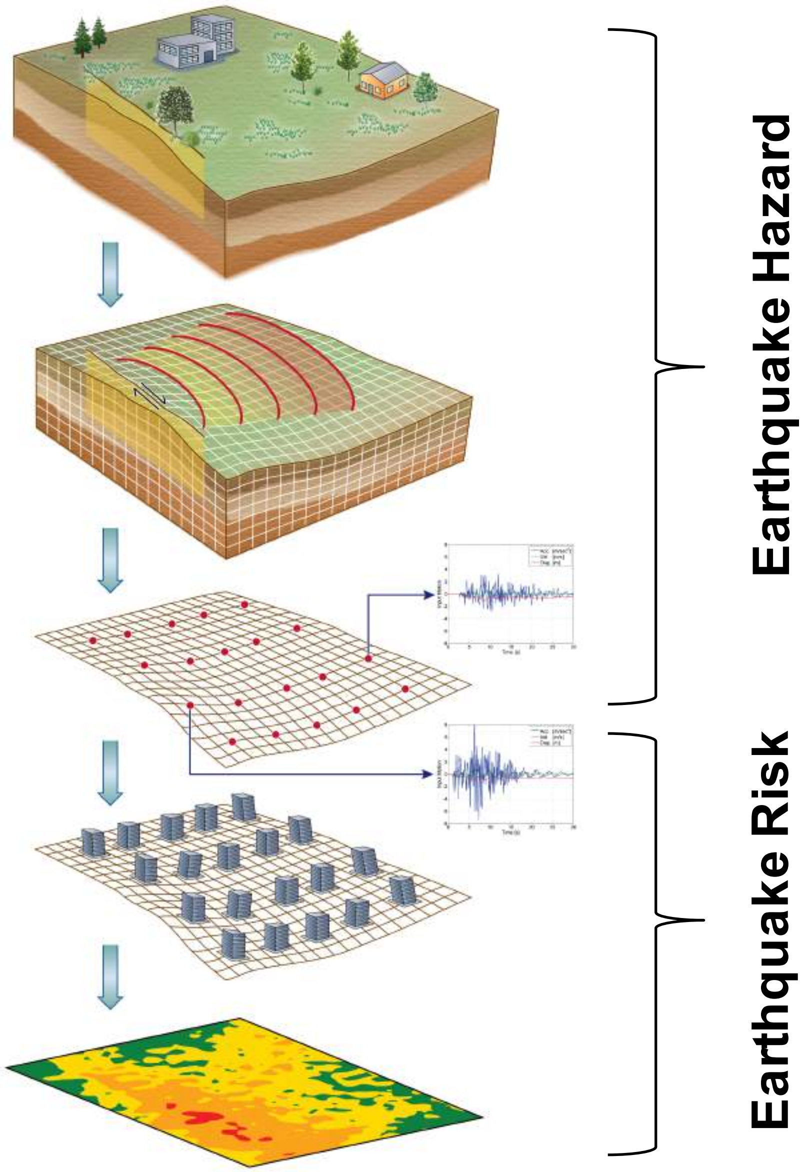

_SW4_ – 4th order finite difference geophysics code for wave propagation

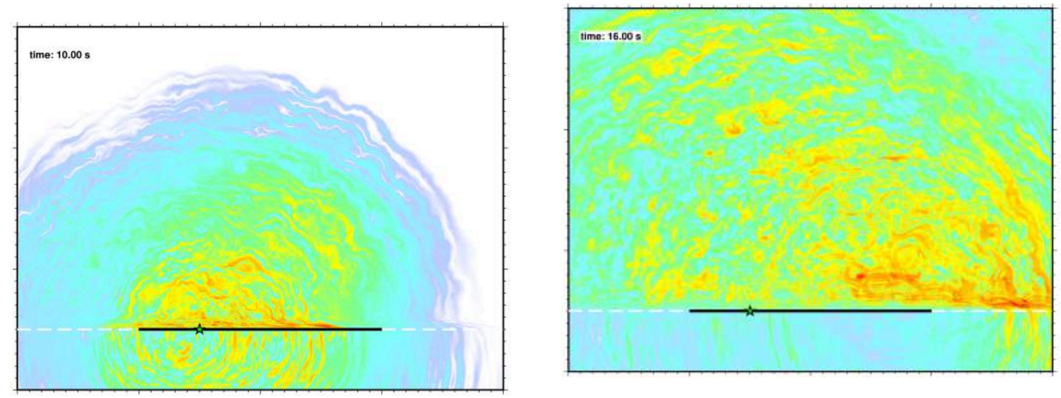

NEVADA & ESSI – finite deformation, inelastic Finite Element codes for structures and soils

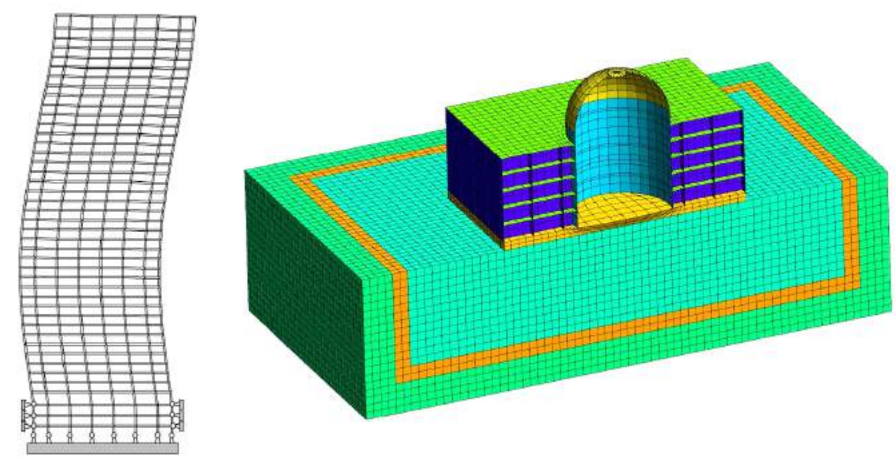

Getting to frequencies of engineering interest is the big computational challenge

Pipelines — **Long-span Bridges** — **Tall Buildings** — **Low-rise Buildings and Industrial Facilities** — **Energy System Components** — **Nuclear Power Equipment**






Exascale objective

0. Hz | **0.1 Hz** | **0.2 Hz** | **1.0 Hz** | **2.0 Hz** | **3.0 Hz** | **10.0 Hz** | **25.0 Hz**

Frequency resolution of ground motion simulations as limited by compute capabilities

Frequency resolution of ground motions simulations as limited by geologic/geotechnical material models

Larger, faster forward simulations

Advanced geologic characterization

Seismologist

Computational challenges to achieving the desired end-state

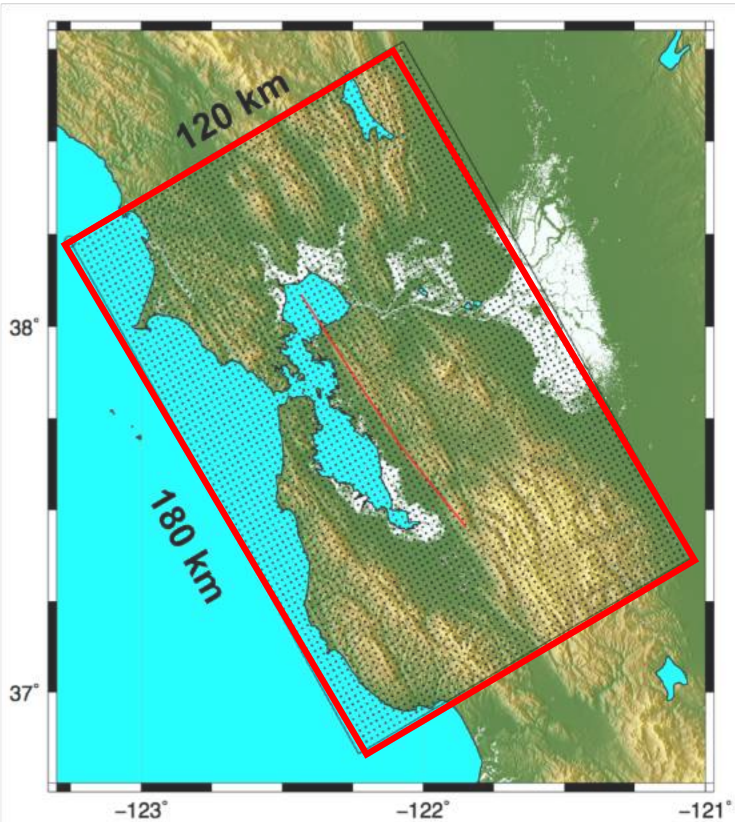

Run much bigger models much faster
- Very large models for resolving 10Hz
- Many realizations to account for uncertainties (e.g. fault rupture)

Representation of fine-scale geology
- Waveform data inversion
- Stochastic geology

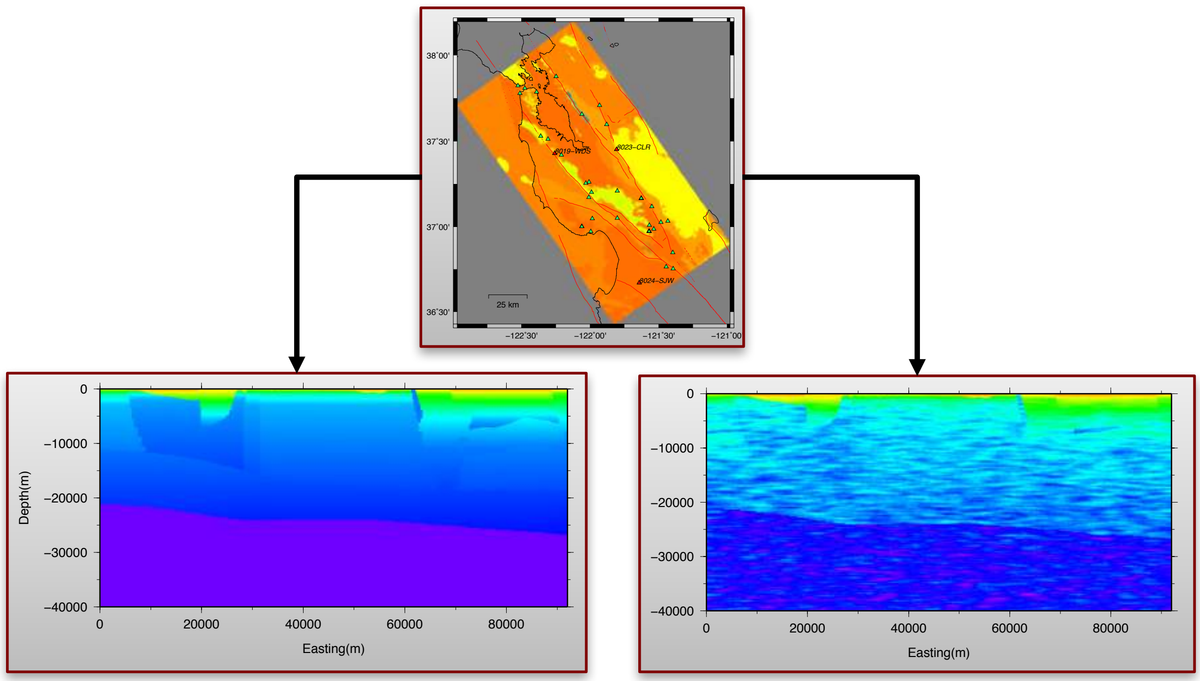

Base geology from data

Base + stochastic geology

A finer point on the objectives

Past

12+ hrs

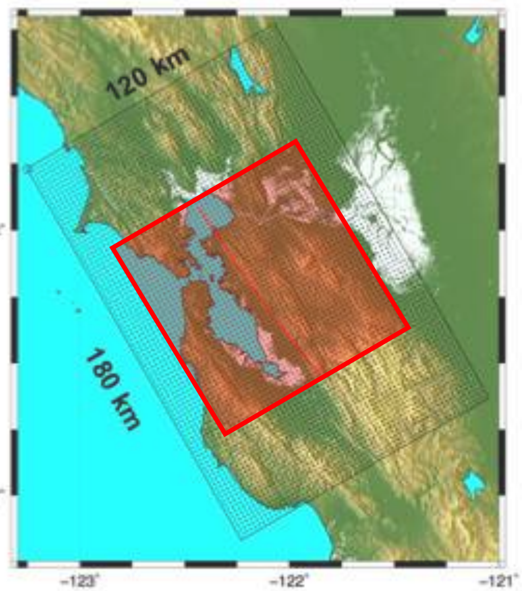

Exascale Future

3-4 hrs?

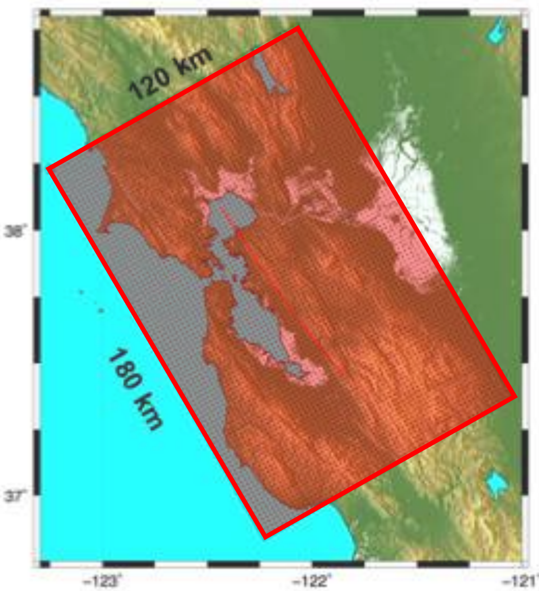

Larger domain

Higher frequency resolution

Simulation times that allow many realizations

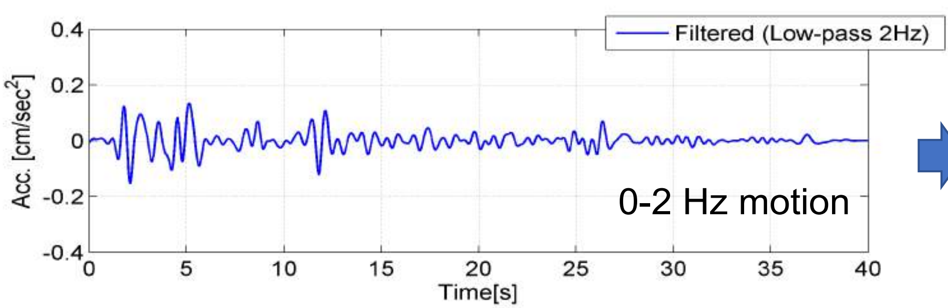

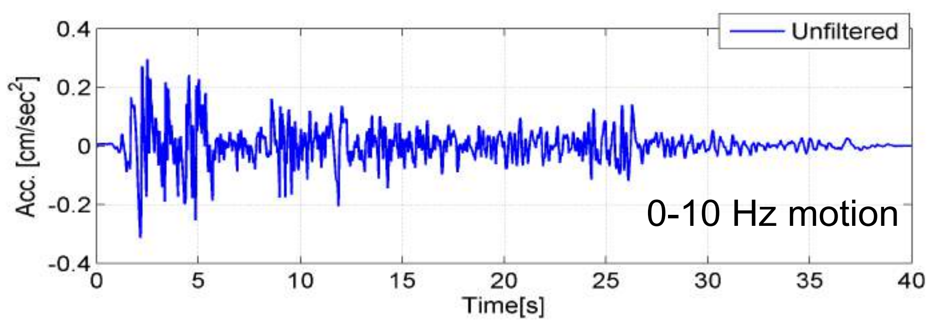

Doubling the frequency resolution = 16X computational effort!

Attributes of SW4 (Seismic Waves 4th Order)

- **SW4 is fourth-order accurate**
  - Explicit time steeping
  - Creates mesh from binary geologic file at run-time
  - Absorbing boundaries etc.
- **Provably stable for…**
  - Heterogeneous materials
  - Creates mesh from binary geologic file at run-time
- **Horizontal MPI task decomposition**
  - Pencil shapes subdomains
  - Easy load balancing
- **Finer meshes and more cores**
  - Pencils get thinner, IO goes up, MPI slows
  - Need to optimize on machine architecture

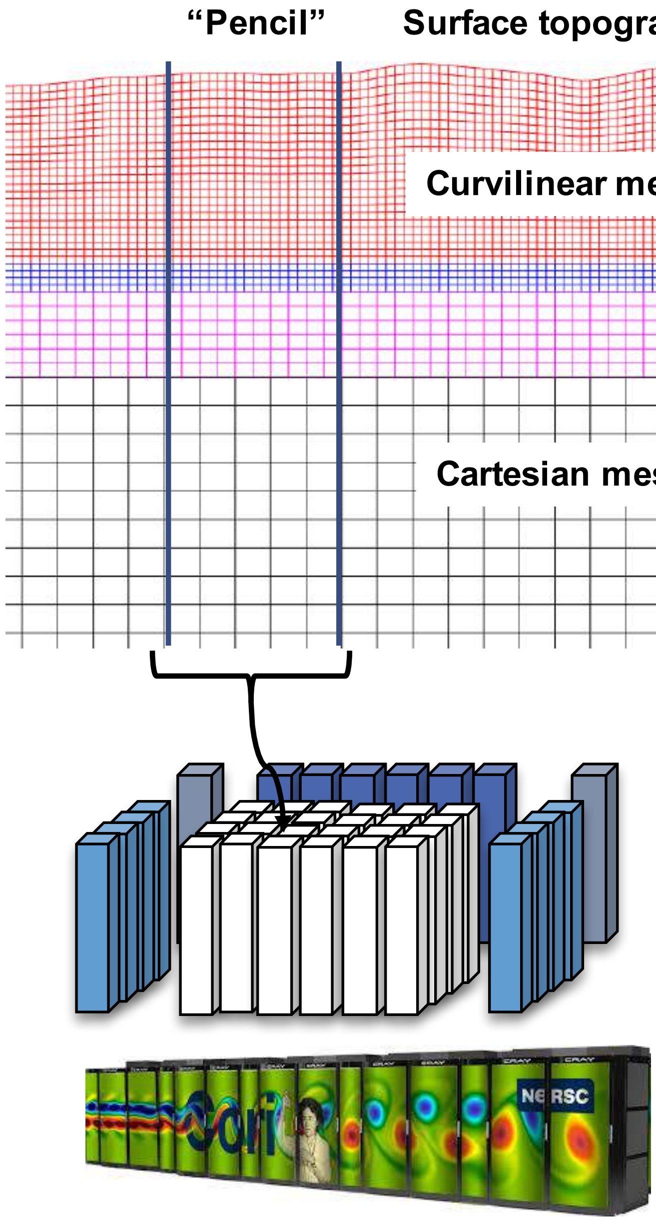

Establishing Key Performance Parameters – development of a suite of test problems

These run fast (~10 min max)

Full hazard-to-risk simulation

Group 1

Group 2

Group 3

Science Demonstration

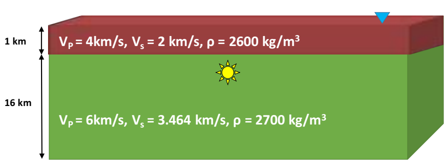

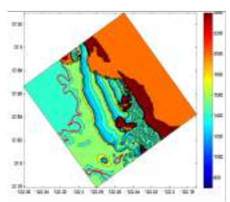

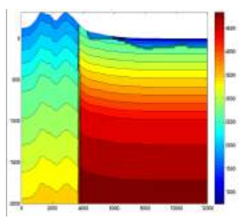

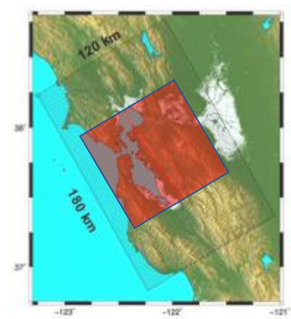

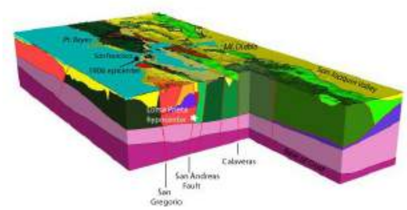

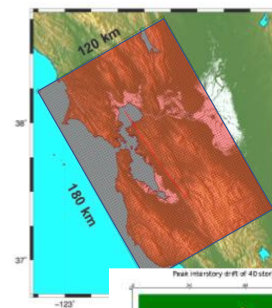

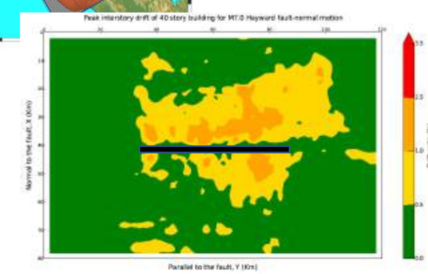

Increasing features/complexity

Test problems for Cartesian mesh refinement FY2018 Q1

Test problems for Curvilinear mesh refinement FY2018 Q2

Test problem for regional-scale simulation performance FY2018 Q3

KPP test problem for regional-scale simulation frequency demonstration FY2018 Q3

A hybrid MPI + OpenMP approach is being utilized for SW4 on CORI

- **MPI supports the message-passing parallel programming model**
  - Explicit calls for passing messages etc.
  - Originally designed for distributed memory architectures
  - Now: distributed, shared or both
- **OpenMP uses threads for parallelism**
  - Shared memory architecture
  - Compiler directives: #pragma omp…
- **SW4 OpenMP within each MPI-task**
  - More grid points per MPI-task
  - Wider computational pencils

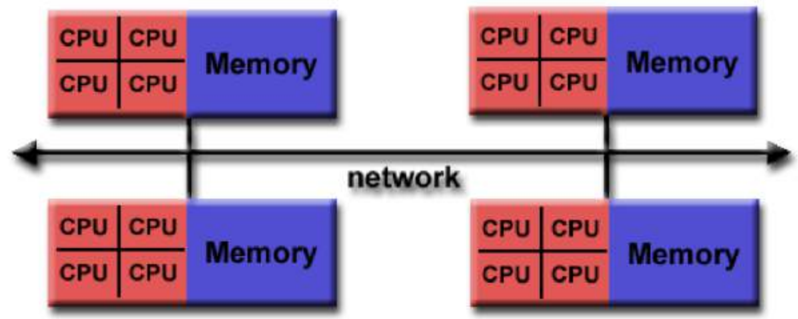

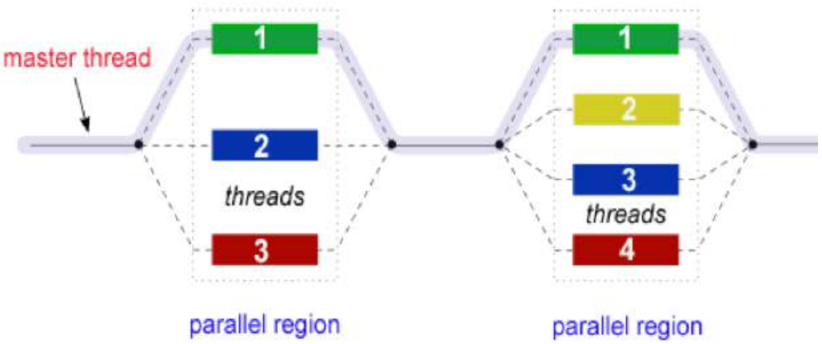

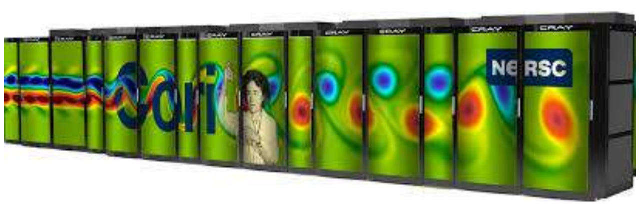

9,668 nodes
68 cores per node

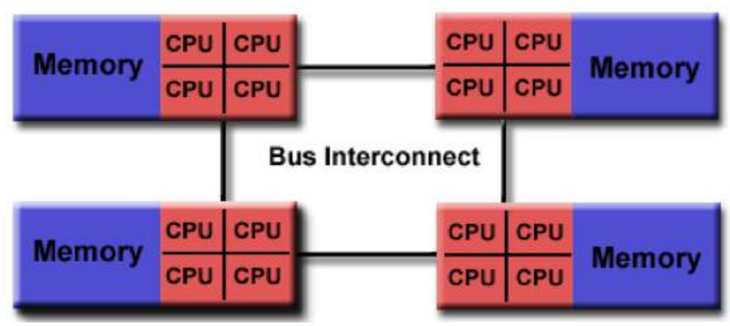

Mesh refinement has been implemented in the SW4 domain (with 4th order accuracy!)

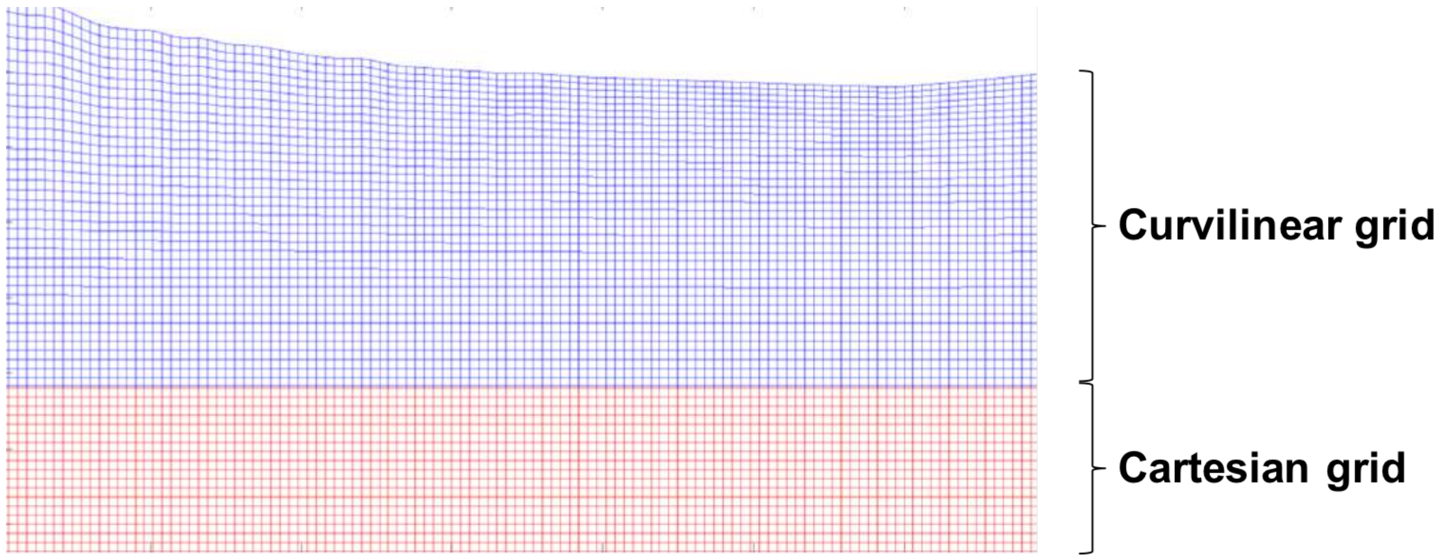

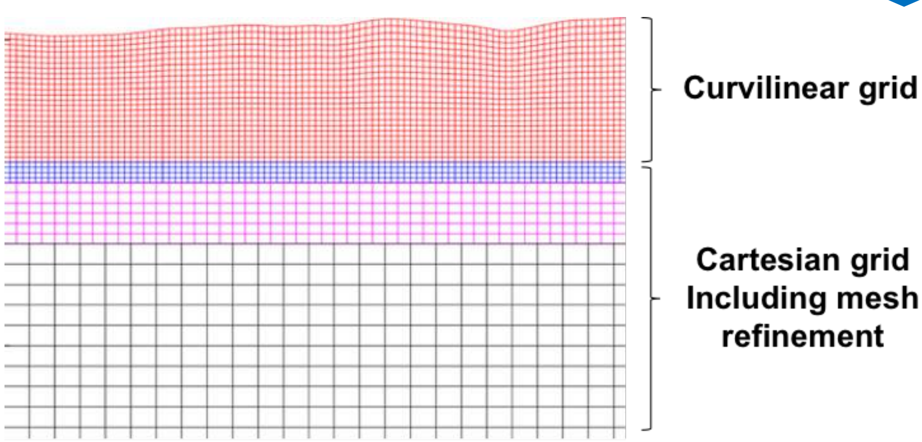

FY18 Q1 milestone
(completed)

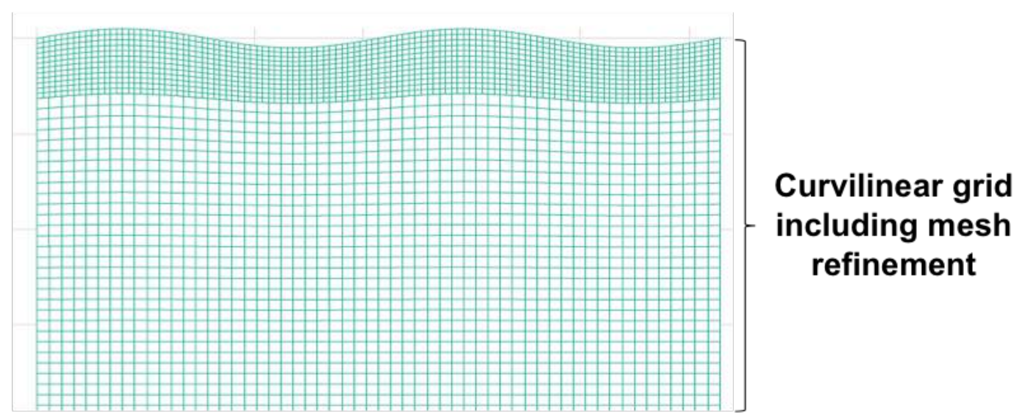

FY18 Q2 milestone
(completed)

We have completed workflow for coupling geophysics and engineering simulations

Earthquake hazard

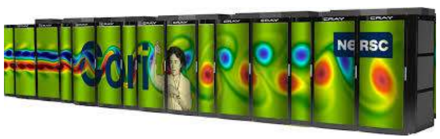

2048 nodes

Surface motions from regional geophysics simulation

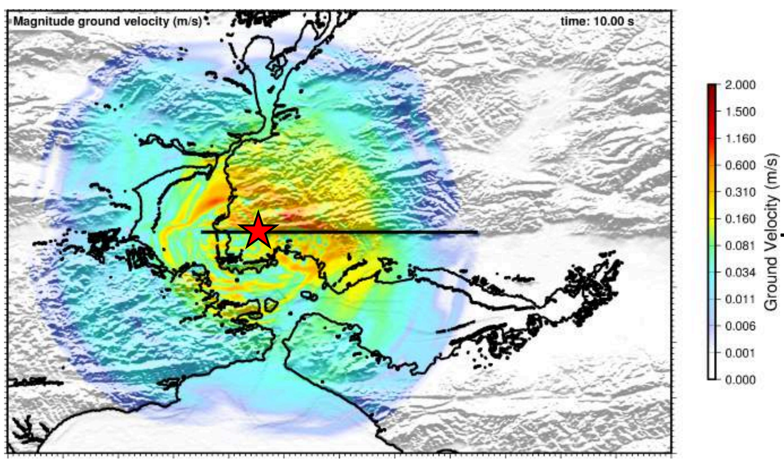

★ Rupture hypocenter

SW4

~ 2000 nonlinear building response history simulations

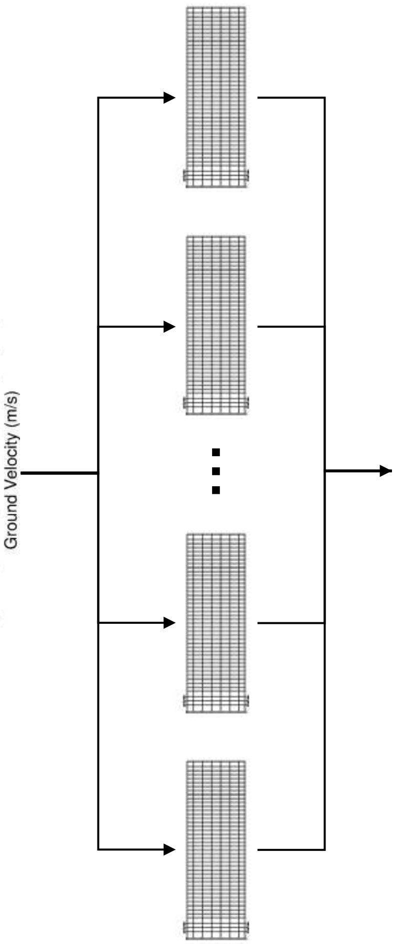

NEVADA

Earthquake risk

50 nodes

Distribution of building peak interstory drift

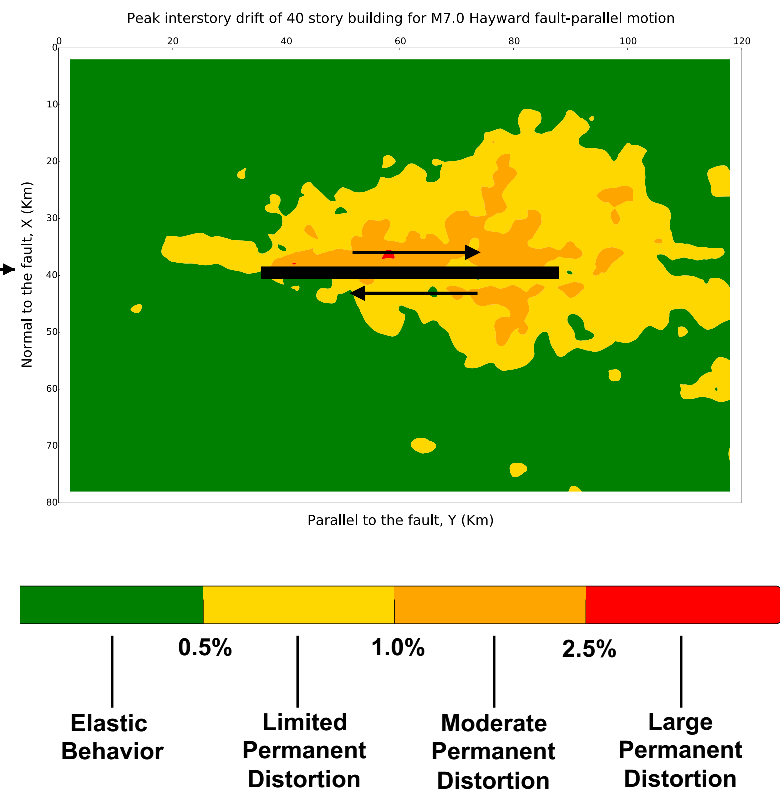

Operational approach

Simulate Earthquake Scenario → **Select Infrastructure Representation** → **Simulate Earthquake Risk**

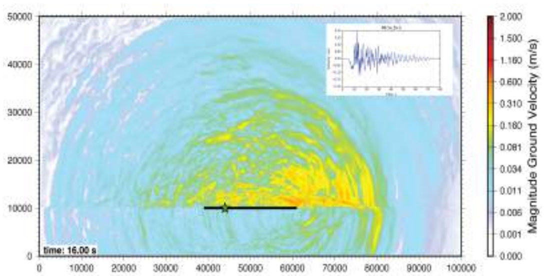

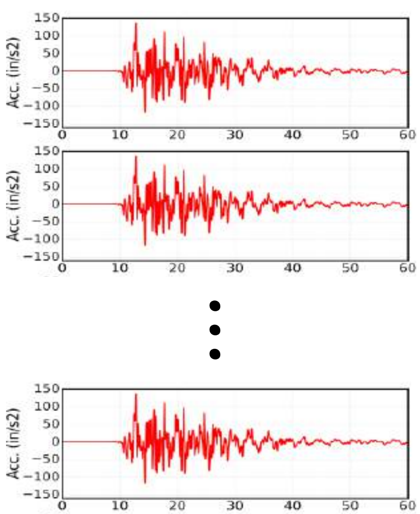

Thousands of ground motions

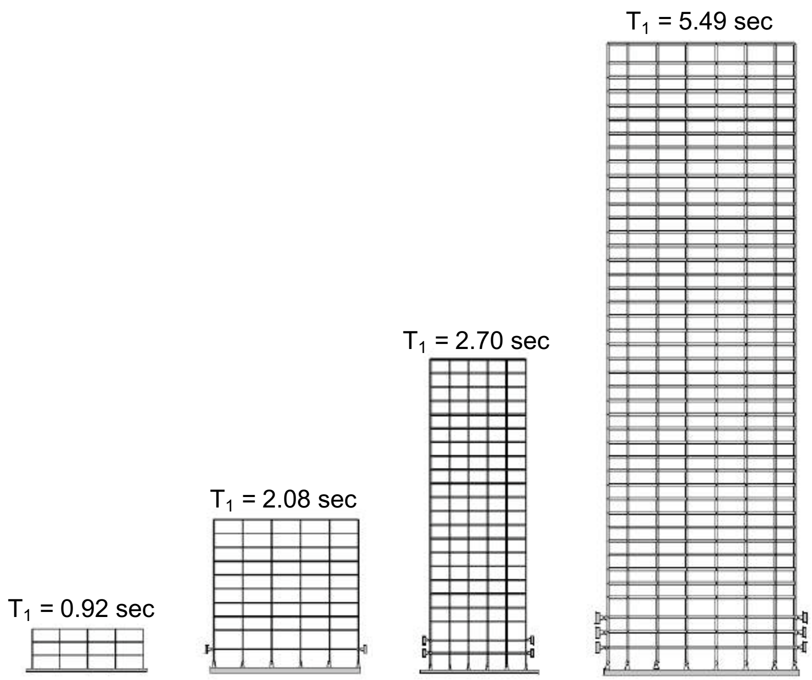

Select infrastructure representation (e.g. nonlinear FEM)

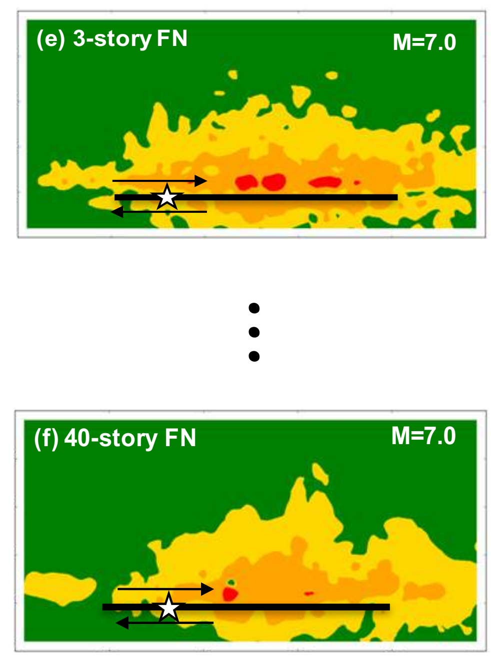

Thousands of response outputs (e.g. peak interstory drift)

In 2017 we completed our first regional scale demonstrations of both hazard and risk

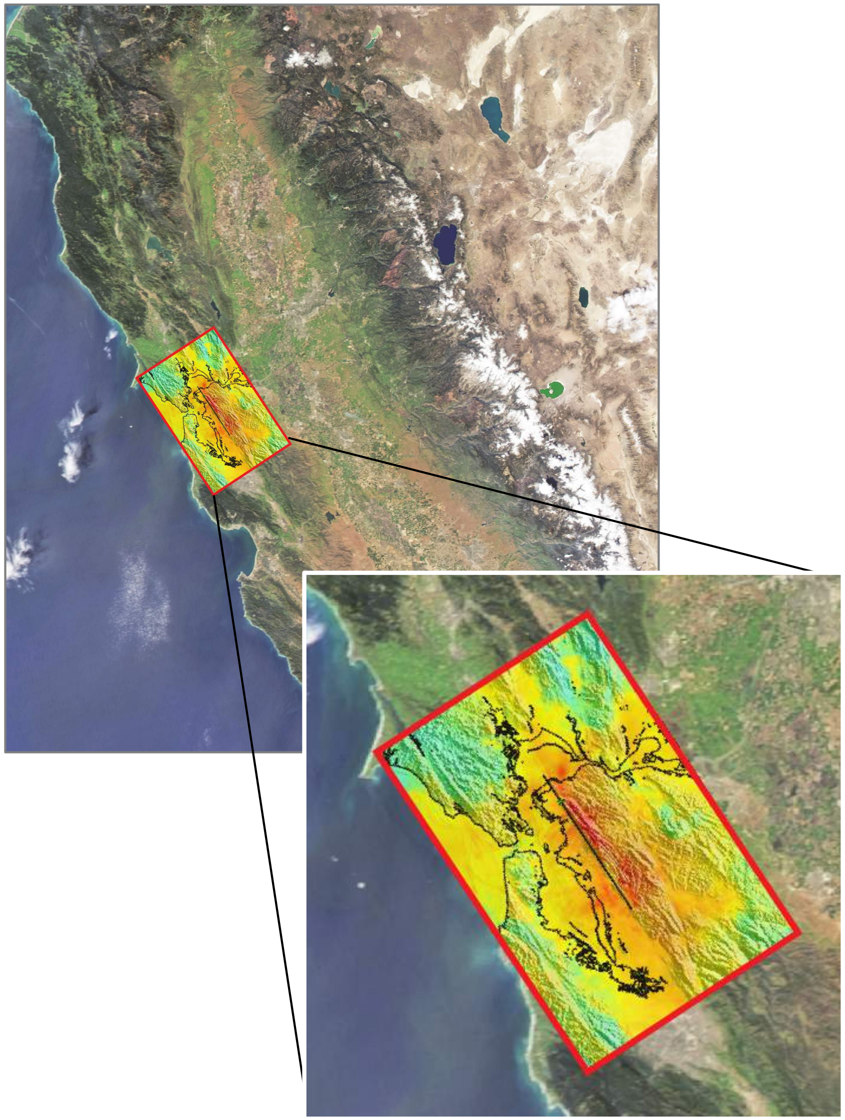

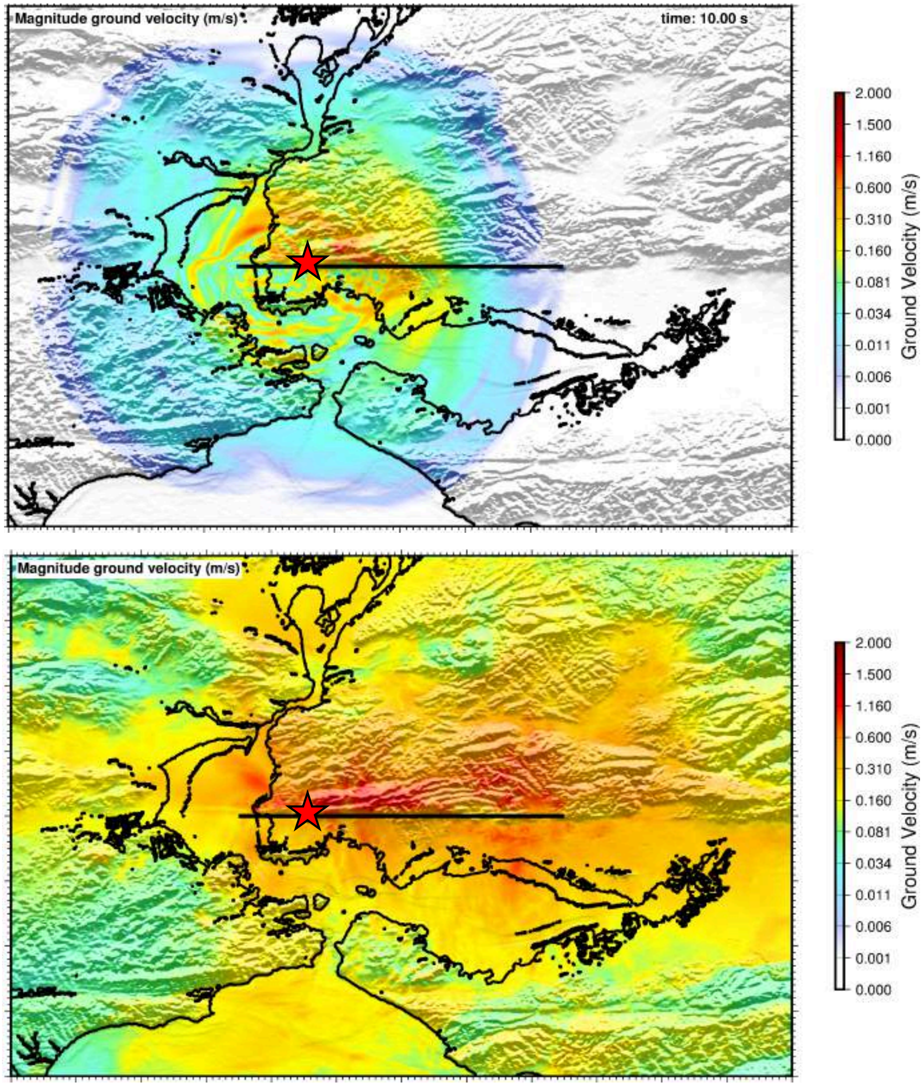

The first regional scale demonstration of simulating both hazard and risk (0-4Hz)

- **Simulation on CORI Phase II (2017 KPP baseline)**

- 87 billion grid points
- 6,528 nodes (2/3 of CORI)
- 417,792 cores
- 12 hour wall clock time
- 5.0 million core hours

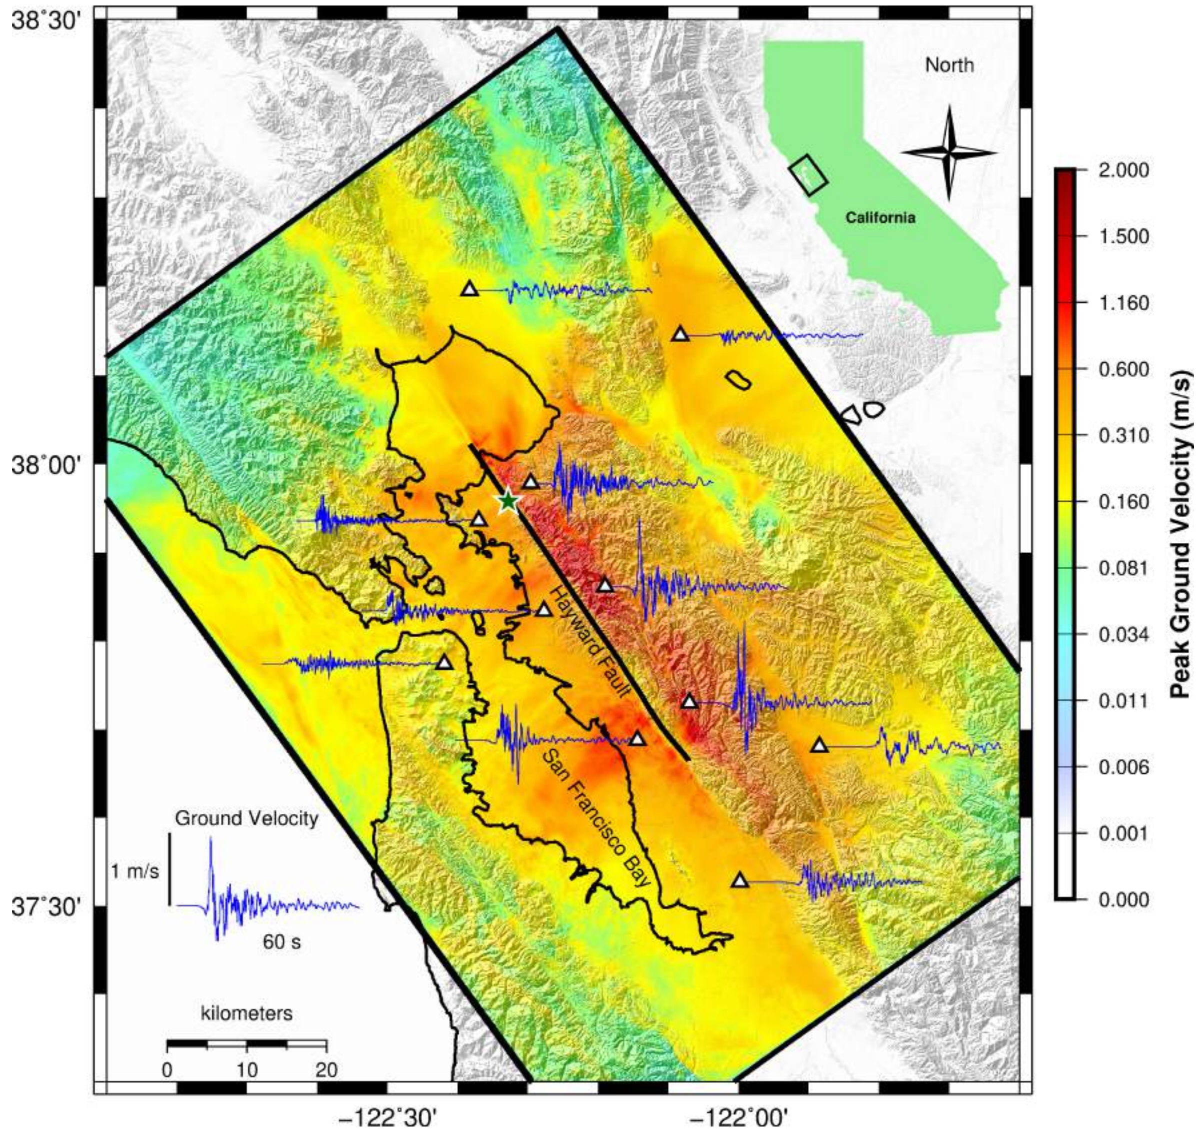

Artie Rodgers et. al. 2018, Seismological Research Letters

M=7 Hayward Fault event

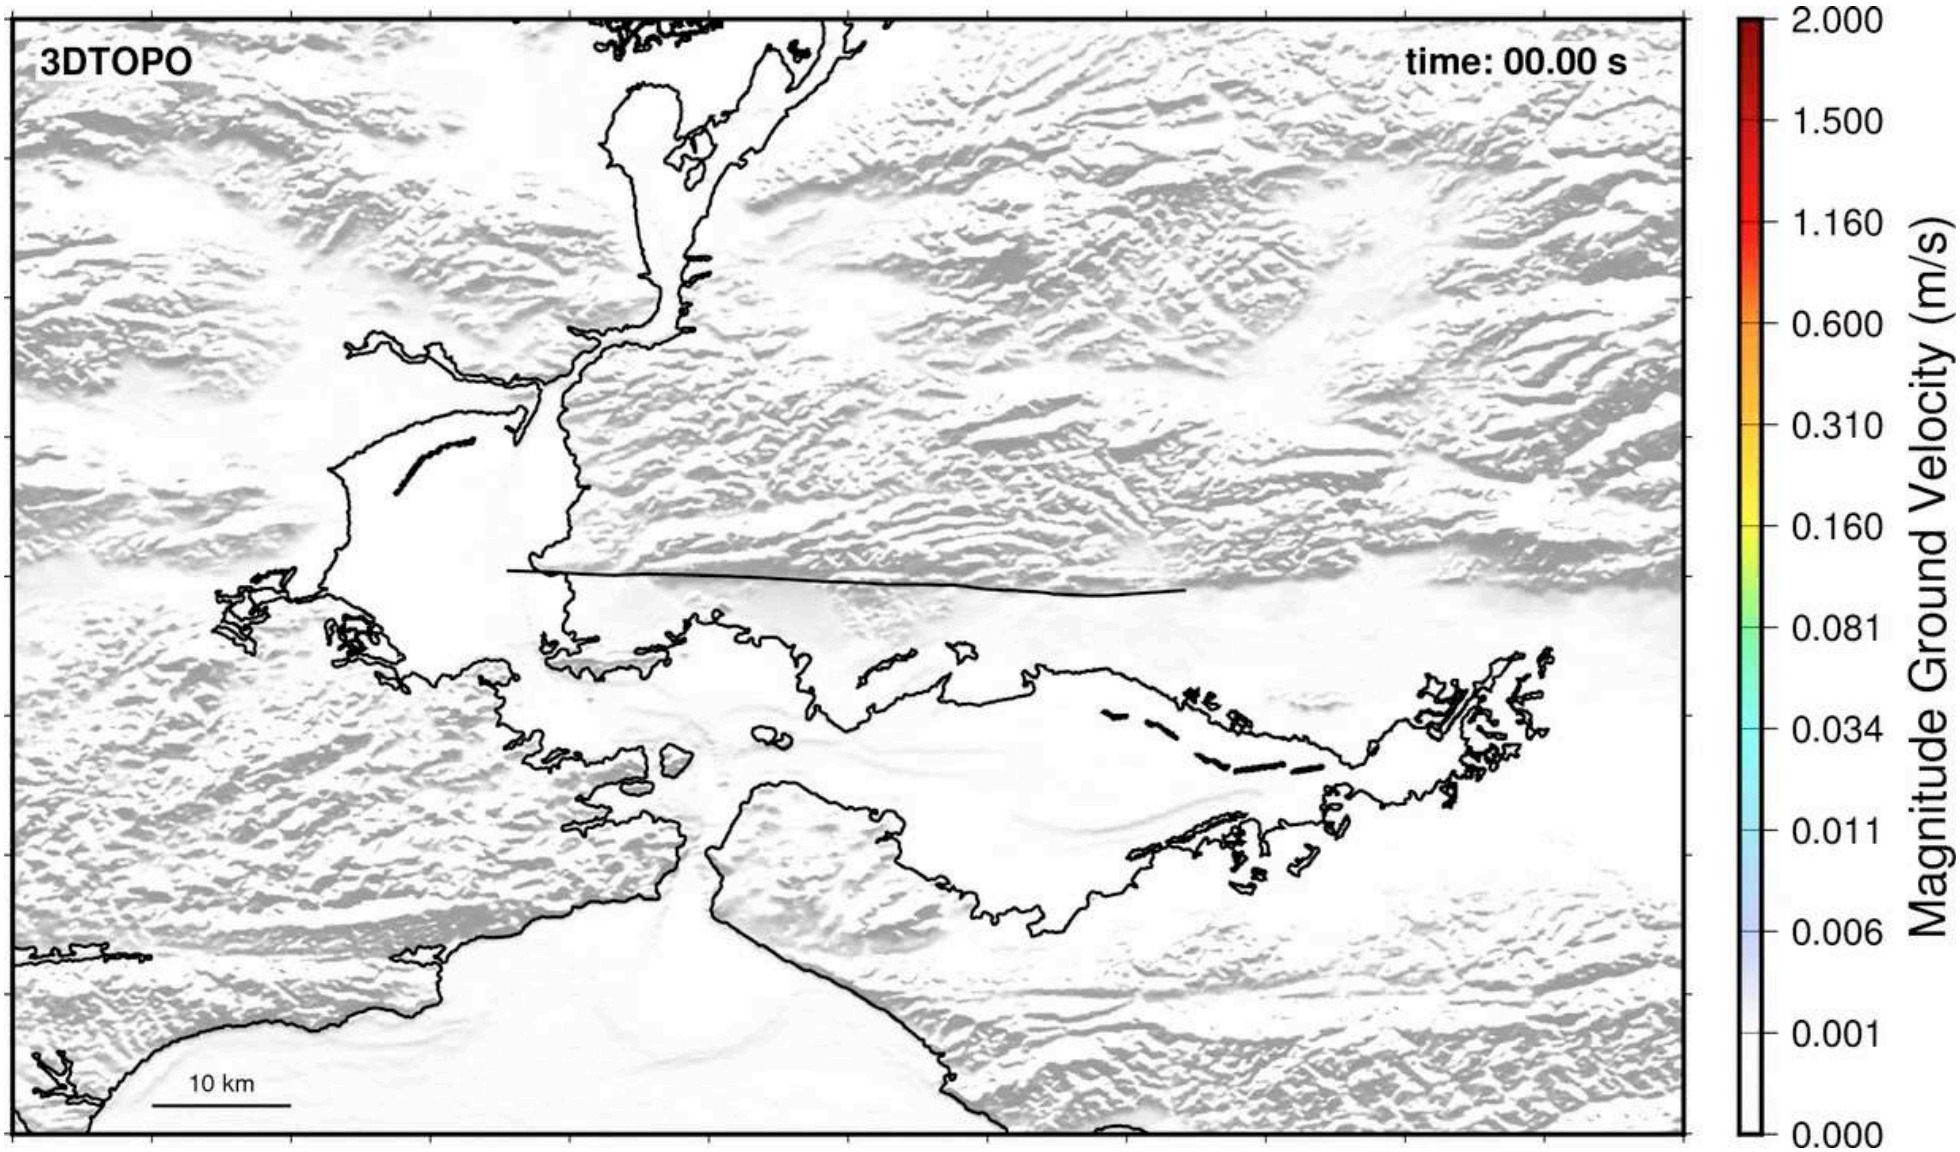

Resulting distribution of risk to three story steel frame buildings (M=7 Hayward event)

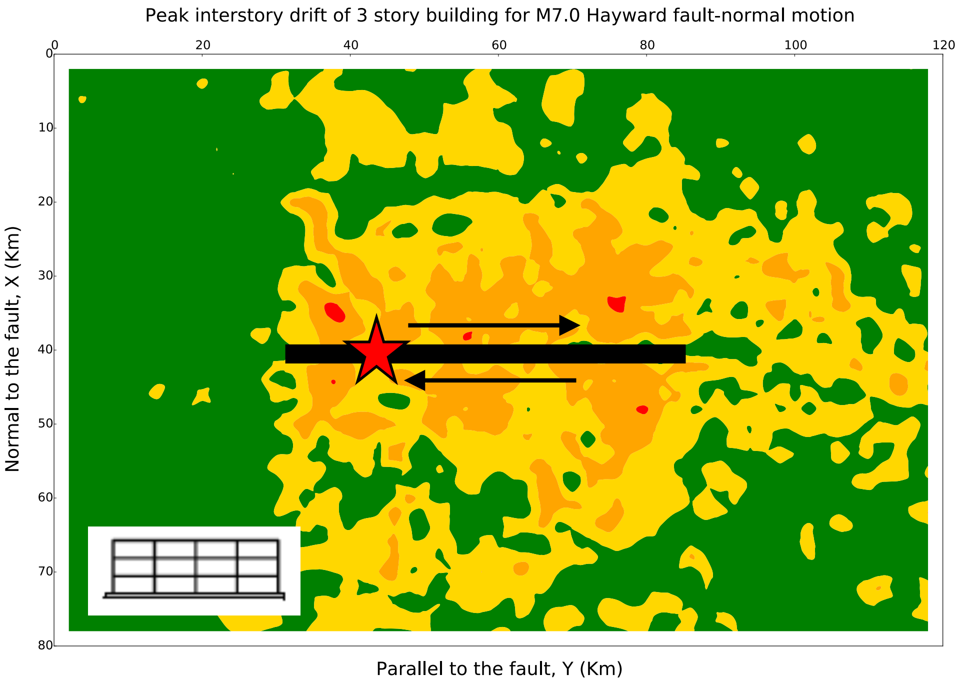

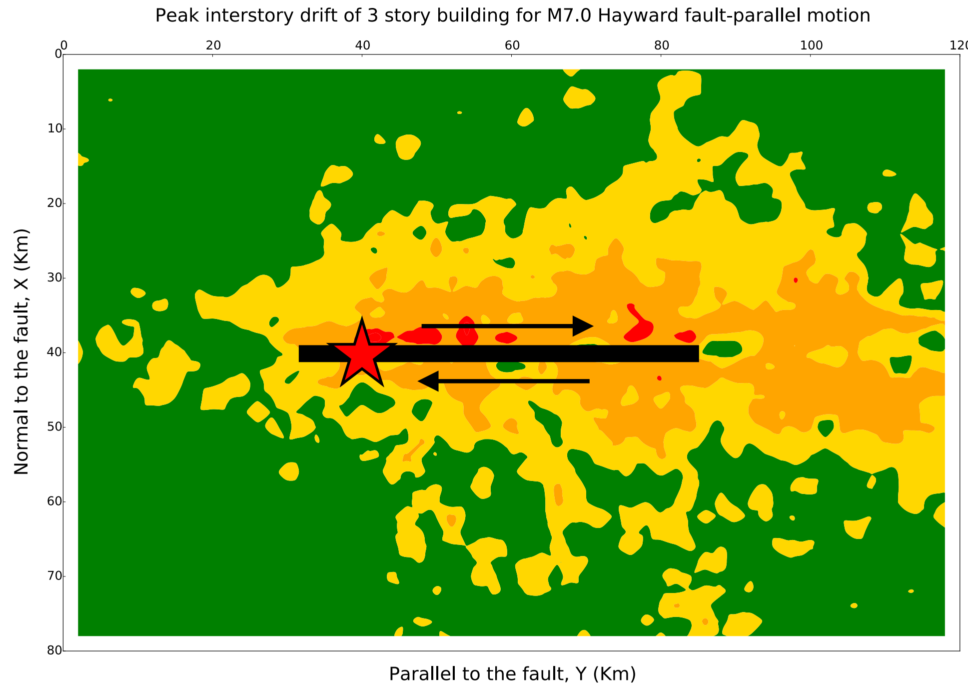

Building Peak Interstory Drift Ratios

0.5% 1.0% 2.5%

Elastic Behavior	Limited Permanent Distortion	Moderate Permanent Distortion	Large Permanent Distortion

DOE standard 1020 limit states

We must critically assess the realism of the simulation results along the way

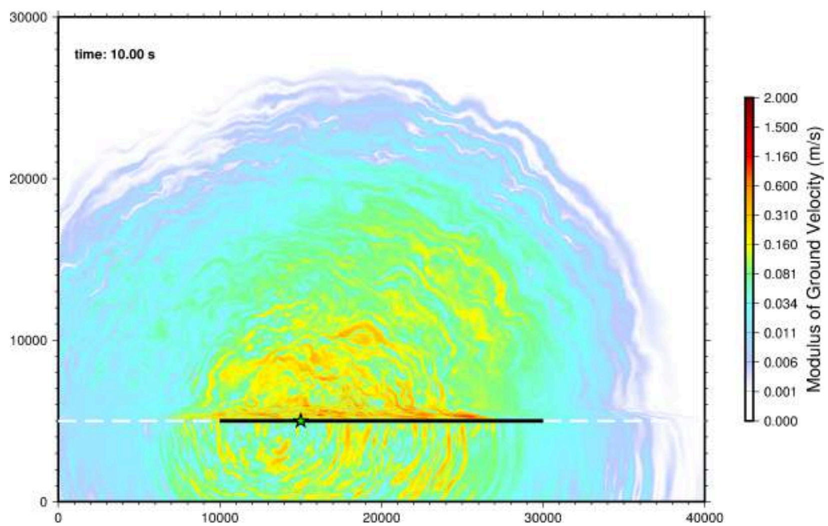

Do ground motion simulations "agree with" observations?

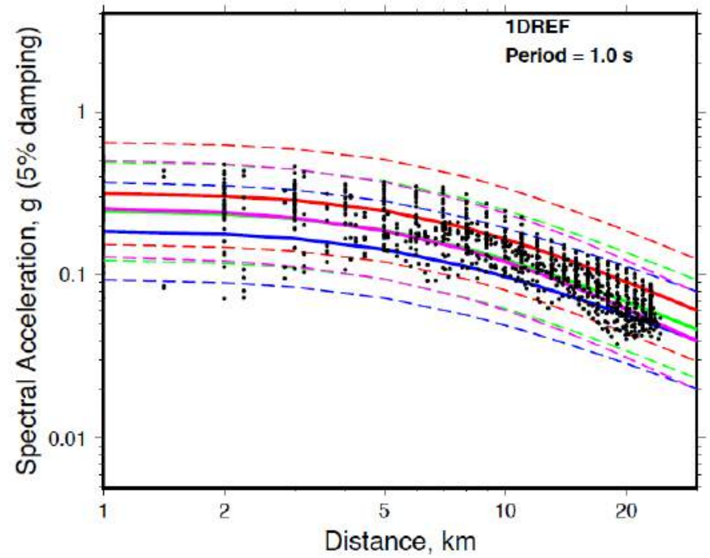

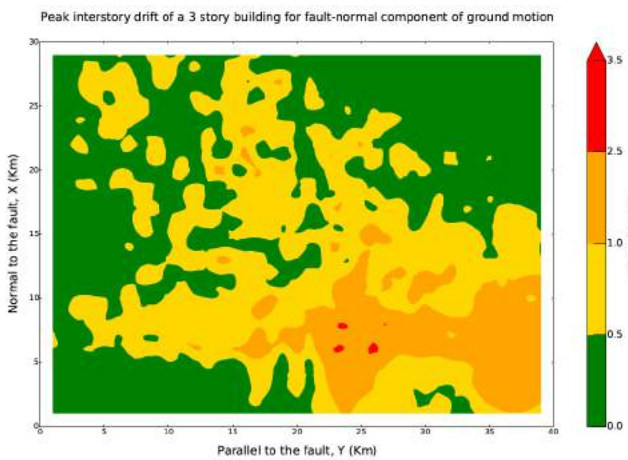

Do structural response simulations "agree with" observations?

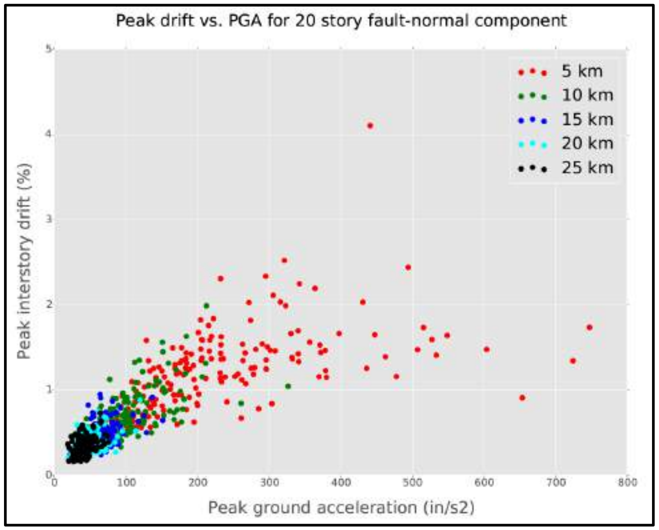

We are seeing some promising things in terms of realism of the ground motions

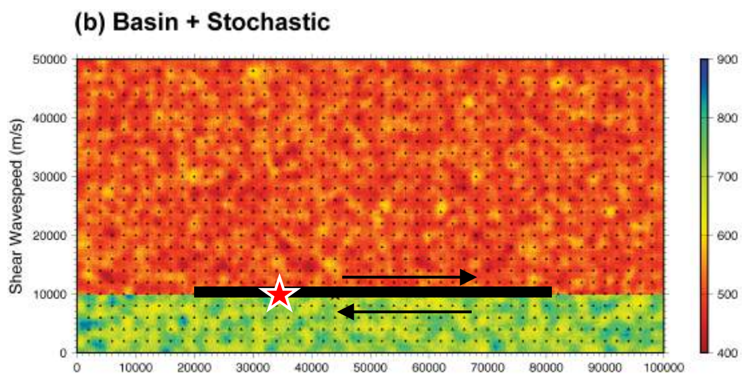

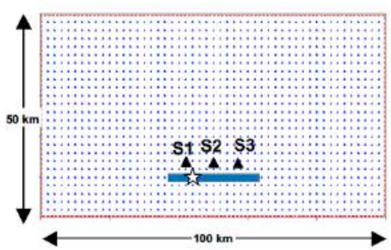

M6.5

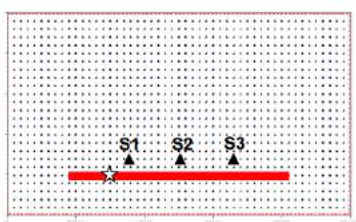

M7.0

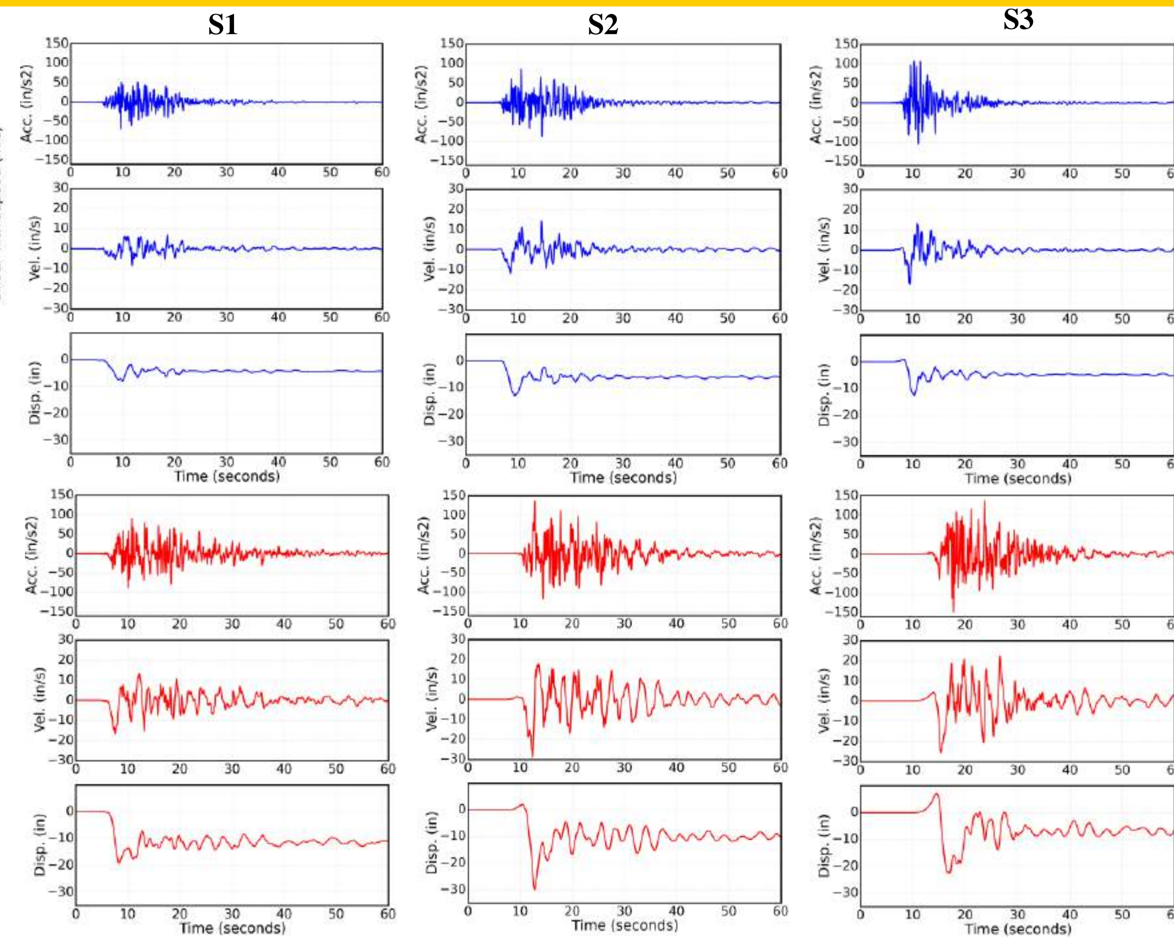

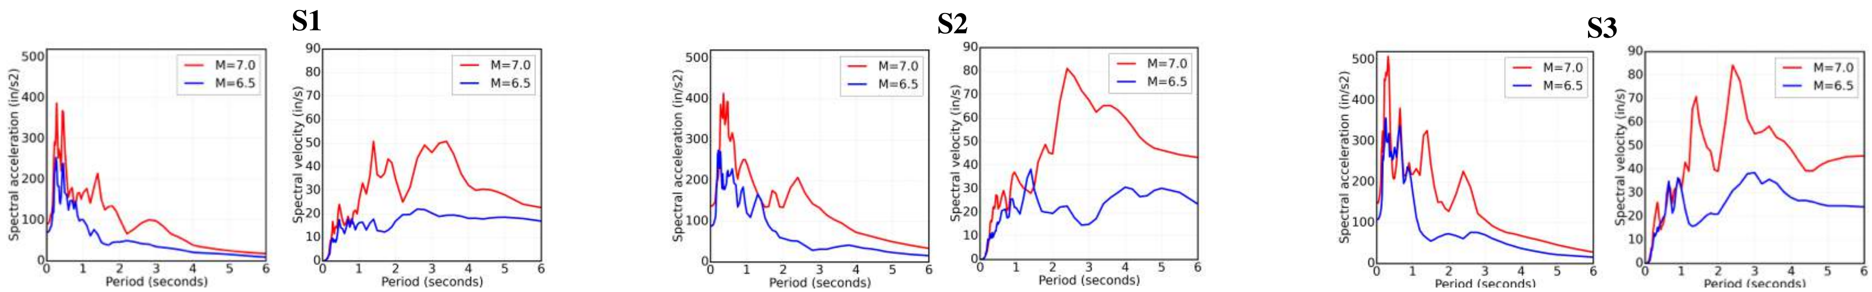

We are seeing some promising things in terms of realism of infrastructure response

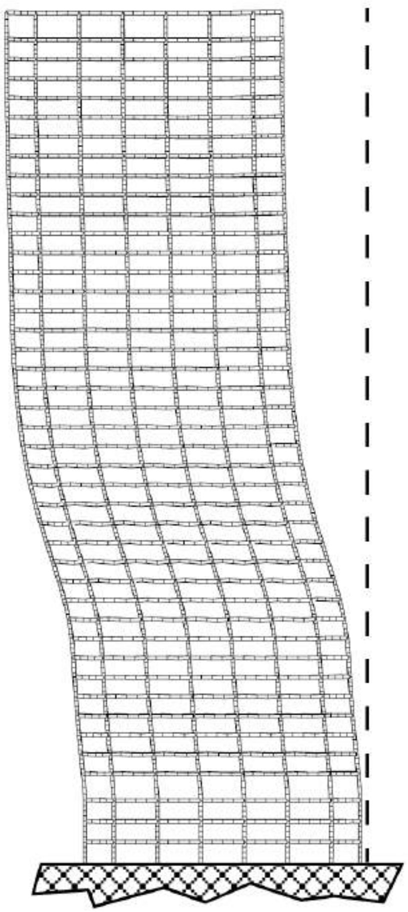

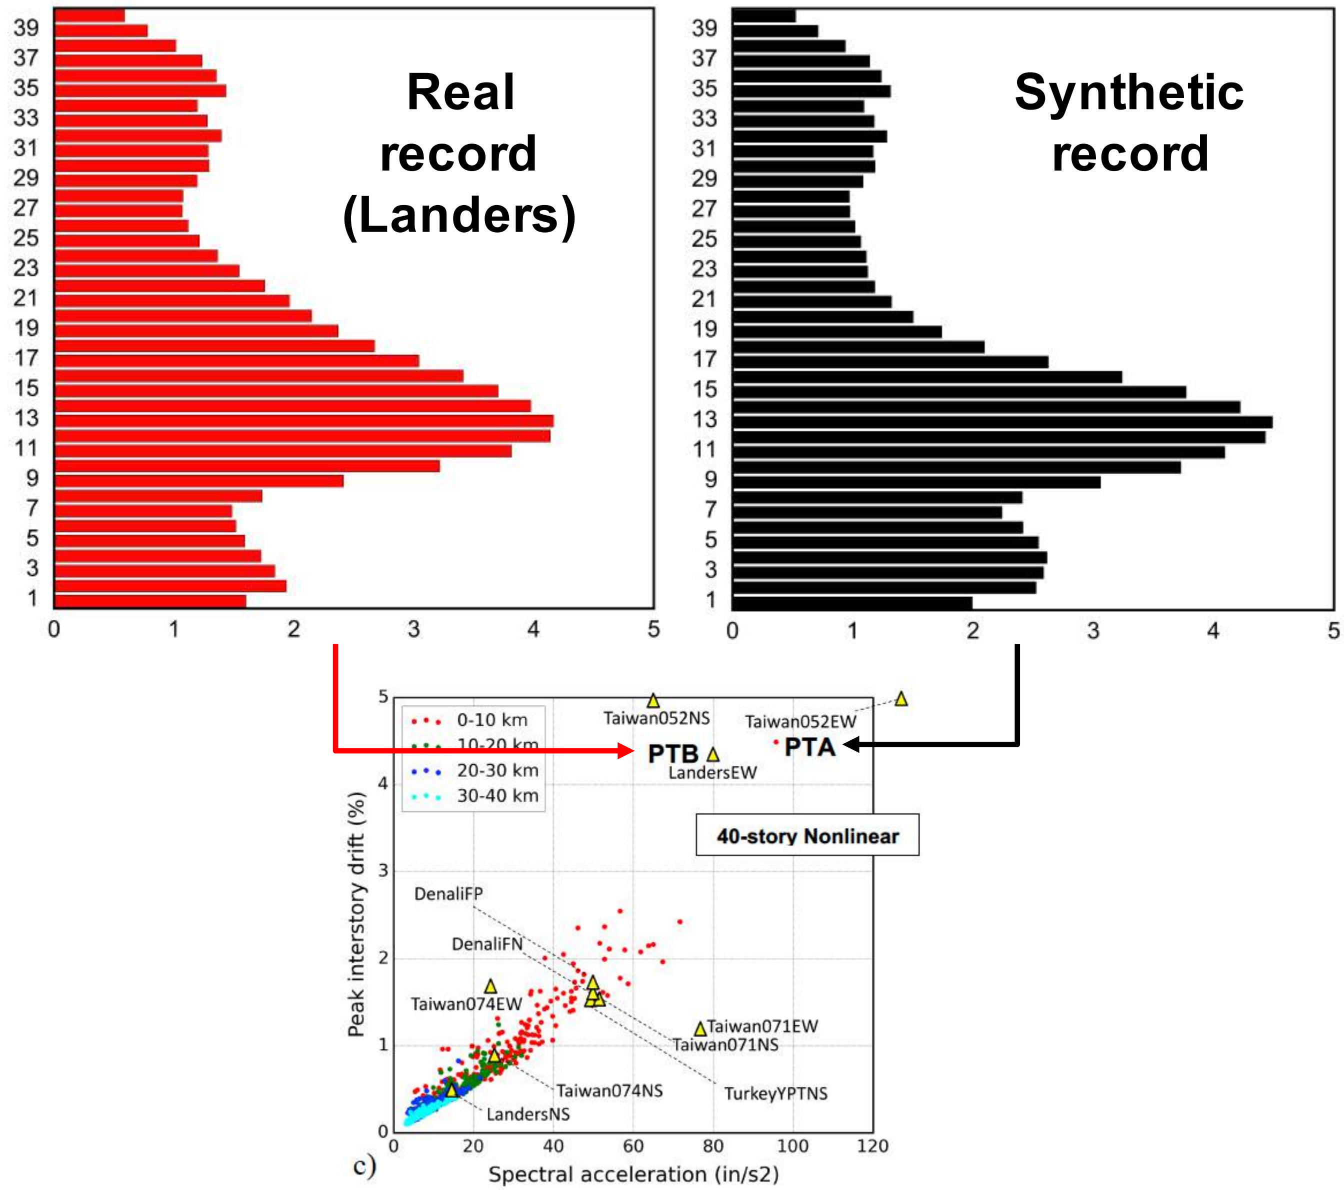

Establishing a Figure of Merit (FOM) for tracking our progress towards ECP goals

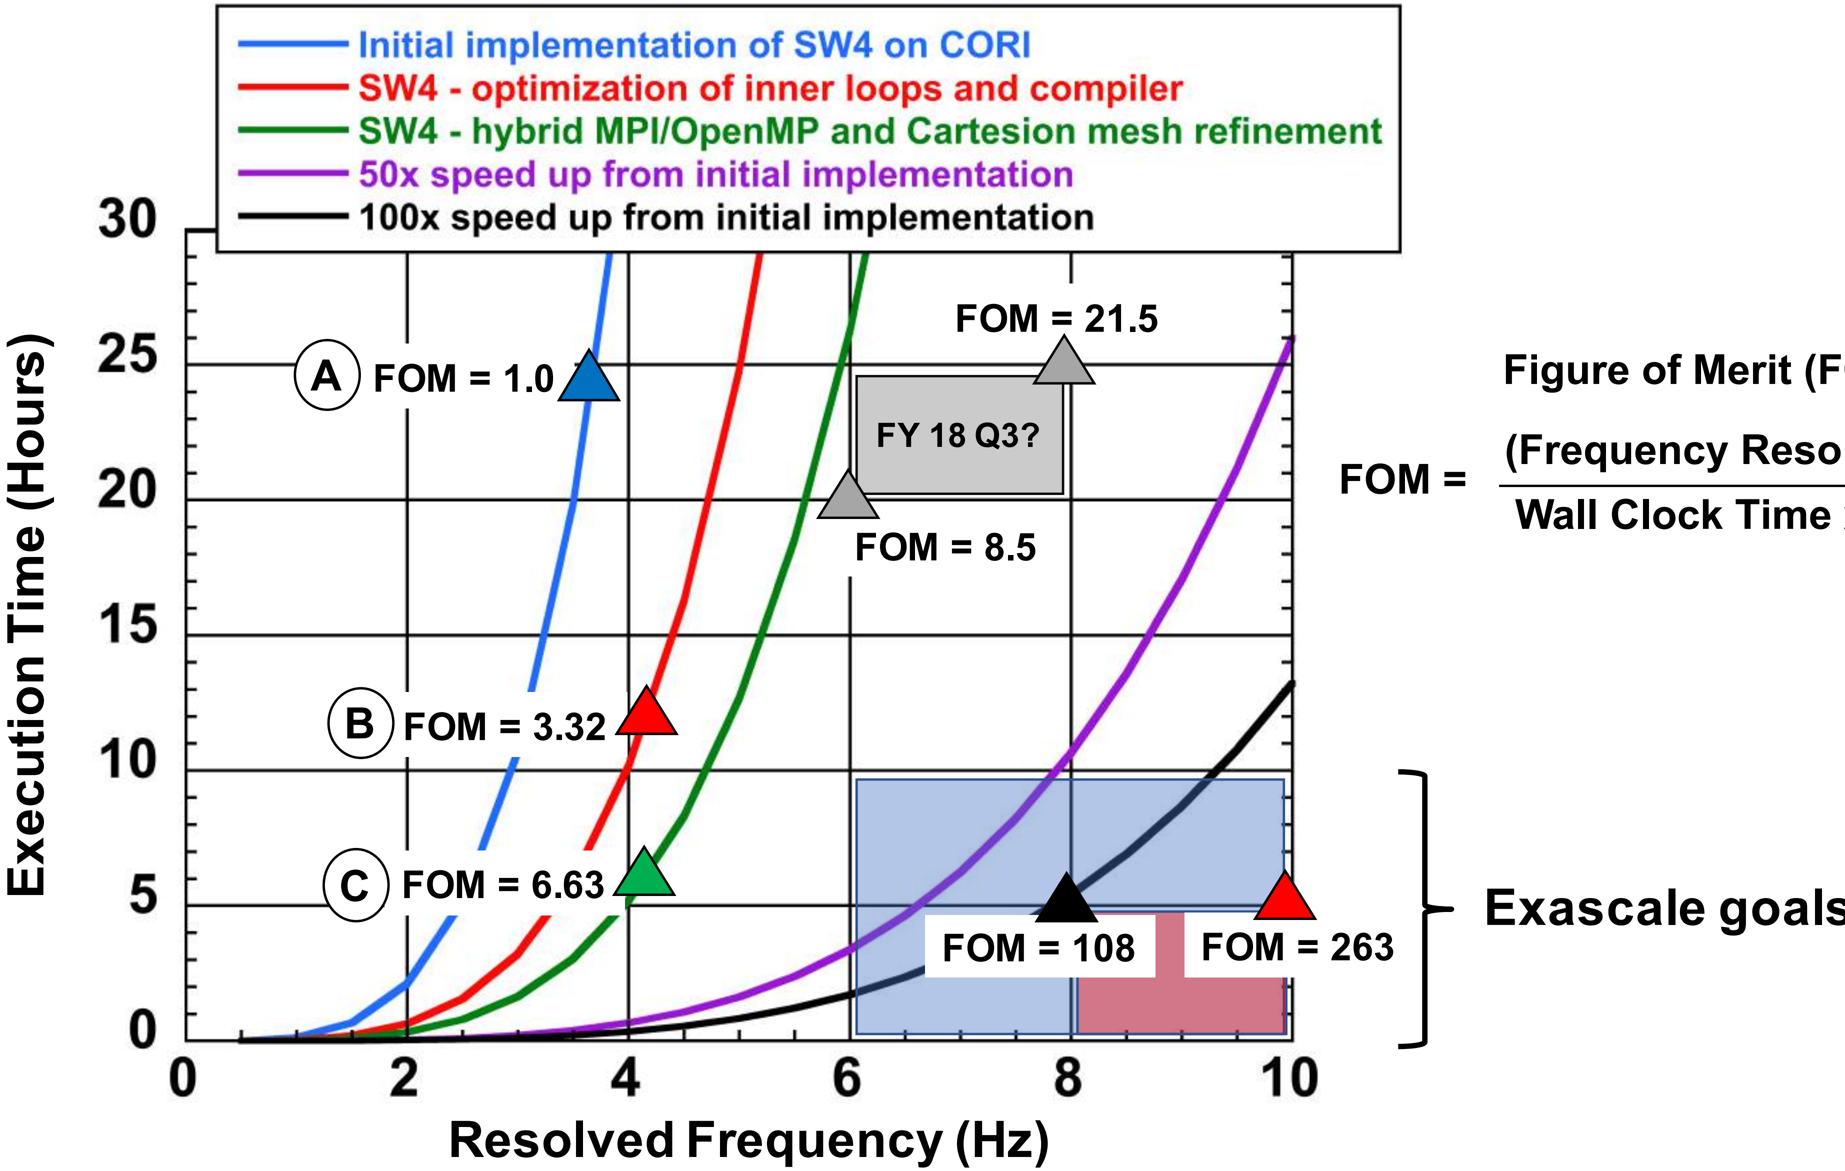

Figure of Merit (FOM)

$$\text{FOM} = \frac{(\text{Frequency Resolved})^4}{\text{Wall Clock Time} \times 7.6}$$

Next on our agenda…

- **Algorithms for waveform inversions**

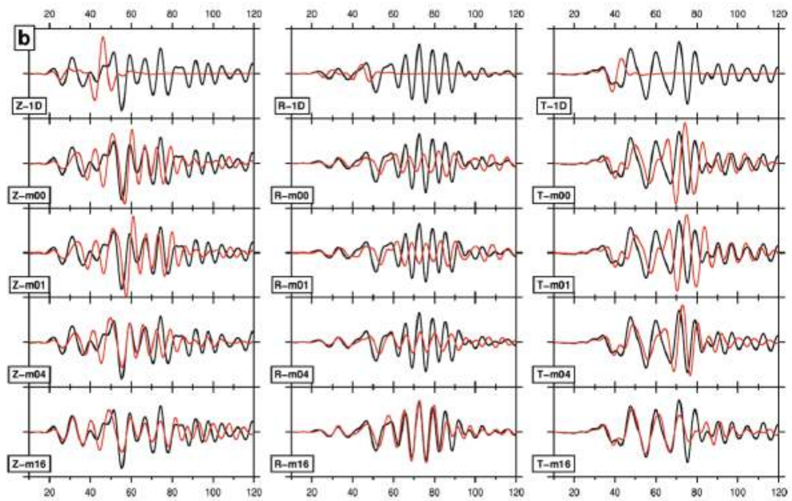

Improved earth structure (FWI)

Ground motion data is becoming available at increased density

- **Preparing for advanced platforms**

porting C++ code to GPU based systems with RAJA C++ libraries

A preliminary assessment has been performed on Ray (4 Nvidia Tesla GPUs per node) and CORI II

RAJA performance on CORI

MPI-OMP	Hand-coded OMP	RAJA OMP
256-1	219.1 sec	238.7 sec
128-2	**216.7 sec**	**247.0 sec**
64-4	259.6 sec	260.4 sec
32-8	226.8 sec	255.9 sec

How far can simulations go, how impactful can they be?

- **Increase our understanding of complex ground motions and interactions between ground motions and structures**

Most certainly – doing this now

- **Augment / improve probabilistic seismic hazard assessments**

Yes – especially understanding path effects

- **Translate to fully simulation-based hazard and risk with appropriate characterization of uncertainties**

Potentially – time will tell

Exascale will give us the tools to find out!