TACC AND NHERI: COMPUTING FOR THE ENDLESS FRONTIER

Dan Stanzione
Executive Director
Associate Vice President for Research

Large Facilities Workshop

April 2019

TACC AT A GLANCE

Personnel
- 175 Staff (~70 PhD)

Facilities
- 12 MW Data center capacity
- Two office buildings, Three Datacenters, two visualization facilities, and a chilling plant.

Systems and Services
- Two Billion compute hours per year
- 5 Billion files, 75 Petabytes of Data, Hundreds of Public Datasets

Capacity & Services
- HPC, HTC, Visualization, Large scale data storage, Cloud computing Consulting, Curation and analysis, Code optimization, Portals and Gateways, Web service APIs, Training and Outreach

A BIT OF HISTORY ABOUT THE TEXAS ADVANCED COMPUTING CENTER (TACC)

TACC LAUNCHED IN JUNE, 2001 AFTER EXTERNAL REVIEW

- Original HPC effort launched by Hans Mark as System initiative in 1986 at $30M
  - Passed between Austin and System several times over next 15 years.
- In 2001, budget of $600k staff of 12 (some shared).
- 50GF computing resource (1/200,000th of the current system).

TOP500 SUPERCOMPUTER SITES

RAPID GROWTH FROM THEN TO NOW…

- 2003 – First Terascale Linux cluster for open science (#26)
- 2004 – NSF funding to join the Teragrid
- 2006 – UT System Partnership to provide Lonestar-3 (#12)
- 2007 - $59M NSF award – largest in UT history – to deploy **Ranger, the world's largest open system (#4)**
- 2008 – funding for new Vis software and launch of revamped visualization lab.
- 2009 - $50M iPlant Collaborative award (largest NSF bioinformatics award) moves a major component to TACC, life sciences group launched.
  - In 2009, we reached, 65 employees.

NOW, A WORLD LEADER IN CYBERINFRASTRUCTURE

- 2010, TACC becomes a core partner (1 of 4) in XSEDE, the TeraGrid Replacement
- 2012, Stampede replaces Ranger with new $51.5M NSF Award
- 2013, iPlant is renewed, expanded to $100M
- 2015, Wrangler, first data intensive supercomputer is deployed.
- 2015, Chameleon cloud is launched
- 2015, DesignSafe, the cyberinfrastructure for natural hazard engineering, is launched.
- 2016 Stampede-2 awarded the largest academic system in the United States, 2017-2021.

HPC DOESN'T LOOK LIKE IT USED TO. . .

HPC-Enabled Jupyter Notebooks
Narrative analytics and exploration environment

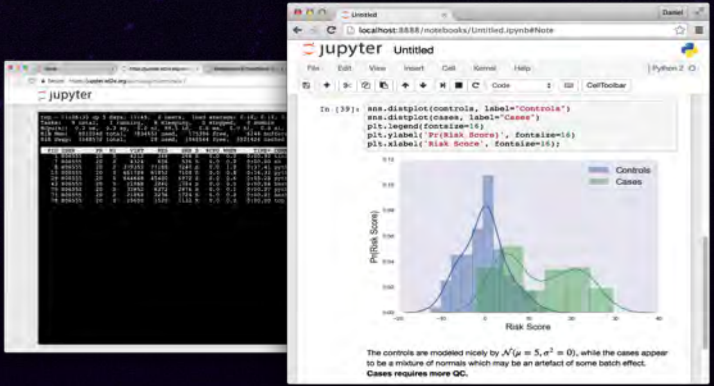

Event-driven Data Processing
Extensible end-to-end framework to integrate planning, experimentation, validation and analytics

Web Portal
Data management and accessible batch computing

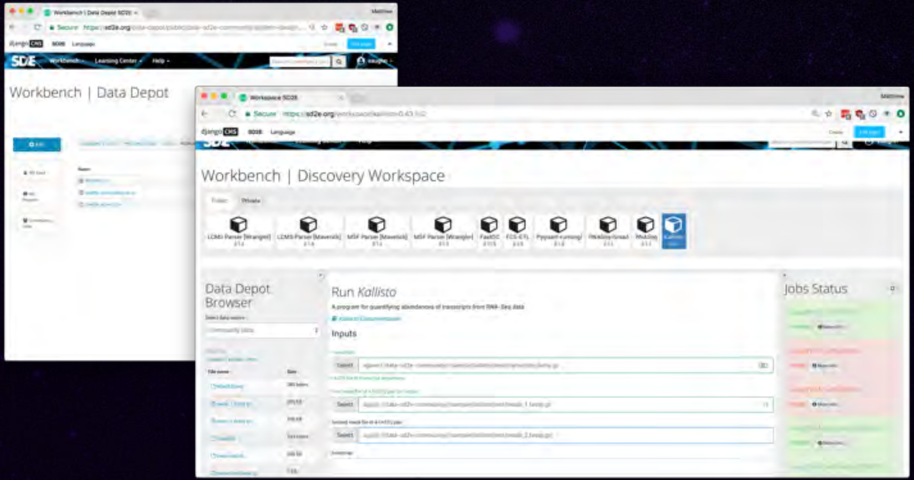

From Batch Processing and single simulations of many MPI Tasks – to that, plus new modes of computing, automated workflows, users who avoid the command line, reproducibility and data reuse, collaboration, end-to-end data management,

- Simulation where we have models
- Machine Learning where we have data or incomplete models

And most things are a blend of most of these. . .

SUPPORTING AN EVOLVING CYBERINFRASTRUCTURE

- Success in Computational/Data Intensive Science and Engineering takes more than systems.
- Modern Cyberinfrastructure requires many modes of computing, many skillsets, and many parts of the scientific workflow.
  - Data lifecycle, reproducibility, sharing and collaboration, event driven processing, APIs, etc.
- Our team and software investments are larger than our system investments
  - Advanced Intefaces – Web front ends, Rest API, Vis/VR/AR
  - Algorithms – Partnerships with ICES @ UT to shape future systems, applications and libraries.

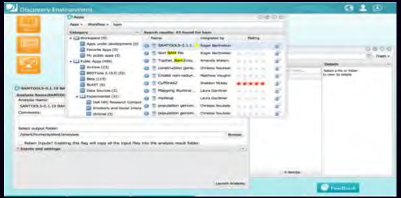

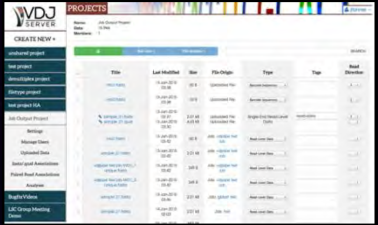

MODERN COMPUTATIONAL SCIENCE

Simulation
Computationally query our
mathematical models of the world

Machine Learning/AI
Computationally query our
data sets
(depending on technique,
also called deep learning)

Analytics
Computationally analyze our
experiments
(driven by instruments that produce
lots of digital information)

I would argue that modern science and engineering combine all three

HIGH PERFORMANCE WORKLOADS

THE THREE CASES WHERE YOU NEED A SUPERCOMPUTER

- TOO BIG (Capability) – Single problems that *won't fit* on a smaller machine.
- TOO MANY (Capacity) – Thousands of small problems that need to be solved to answer a question, that would take too long on a smaller machine
- REAL TIME – Hard deadline for completion, i.e. hurricane forecasting

AN ECOSYSTEM FOR EXTREME SCALE SUPERCOMPUTING

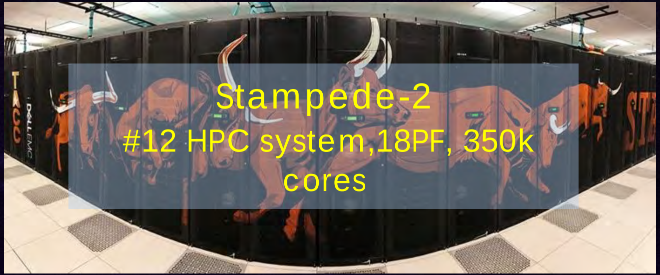

Stampede-2
#12 HPC system,18PF, 350k cores

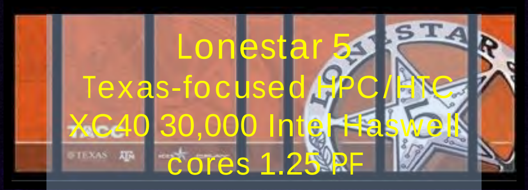

Lonestar 5
Texas-focused HPC/HTC
XC40 30,000 Intel Haswell cores 1.25 PF

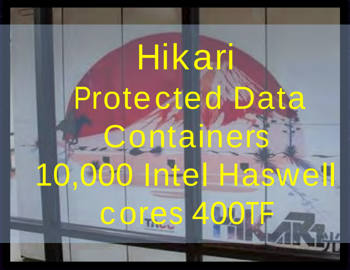

Hikari
Protected Data Containers
10,000 Intel Haswell cores 400TF

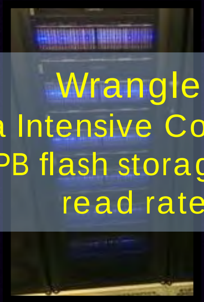

Wrangler
Data Intensive Computing
0.6 PB flash storage 1 TB/s read rate

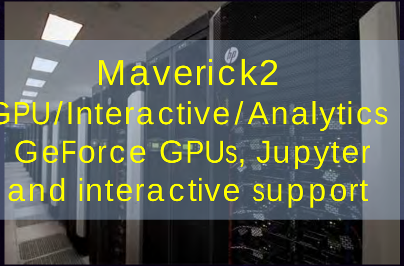

Maverick2
GPU/Interactive/Analytics
GeForce GPUs, Jupyter and interactive support

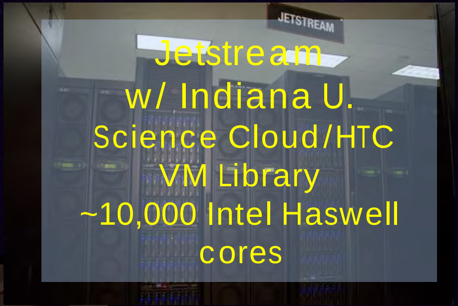

Jetstream
w/ Indiana U.
Science Cloud/HTC
VM Library
~10,000 Intel Haswell cores

Stockyard
Shared Storage Across TACC
30PB, Lustre

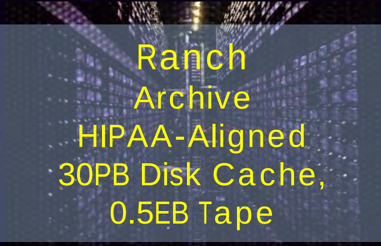

Ranch
Archive
HIPAA-Aligned
30PB Disk Cache,
0.5EB Tape

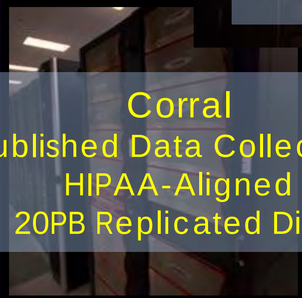

Corral
Published Data Collections
HIPAA-Aligned
20PB Replicated Disk,

Rodeo

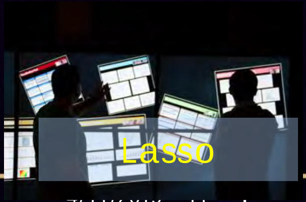

Lasso

EXPERIMENTAL SYSTEMS

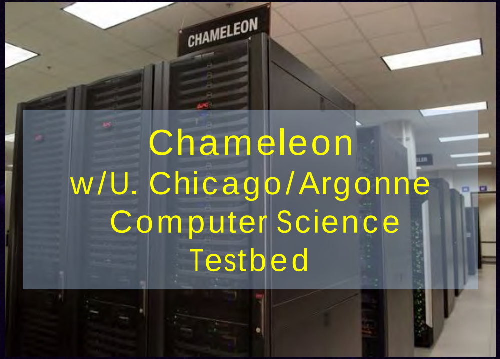

Chameleon
w/U. Chicago/Argonne
Computer Science
Testbed

Catapult
Altera FPGA Testbed
(Microsoft)

Fabric
Alternate Architectures
(IBM,CAPI,FPGA, GPU)

Discovery
New Processor/Storage
Benchmarking

Rustler
Object Storage Testbed

AN EXEMPLAR PROJECT – SD2E

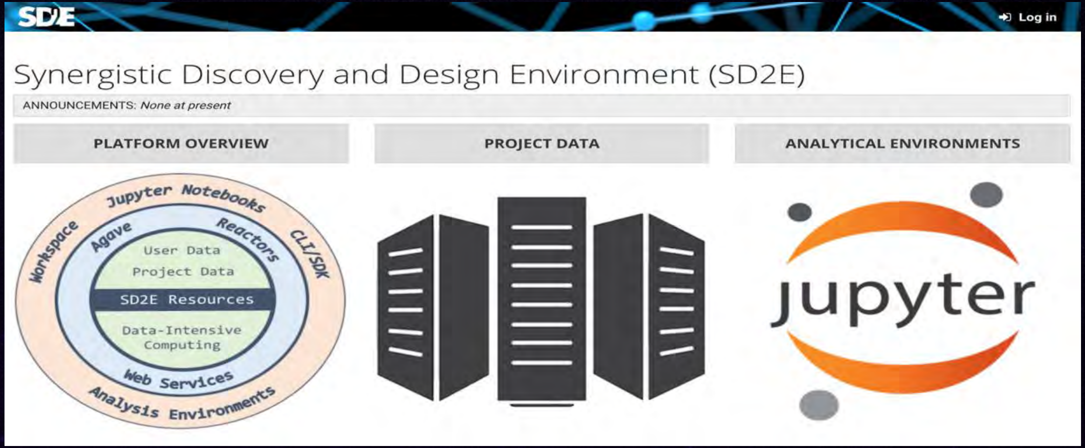

- DARPA – "*Synergistic Discovery and Design* (SD2)"
- Vision: to "develop data-driven methods to accelerate scientific discovery and robust design in domains that lack complete models."
- Initial focus in synthetic biology; ~six data provider teams, ~15 modeling teams, TACC for platform
- Cloud-based tools to collect, integrate, and analyze diverse data types; Promote collaboration and interaction across computational skill levels; Enable a reproducible and explainable research computing lifecycle; Enhance, amplify, and link the capabilities of every SD2 performer

TACC SUPPORTS AN INCREDIBLE AMOUNT & DIVERSITY OF RESEARCH

- Our request rate (NSF Sytems) continues to be about 5-10x what we can deliver
  - More than 2,000 unique users run jobs in any given month (Stampede2)
    - (More than 12k people have run in production on Stampede2, spanning 3,500+ projects).
  - 2+ **million** successful jobs last year.
  - We estimate well 0ver 35,000 use TACC systems via Web or API.

COSMOS GRAVITATIONAL WAVES STUDY

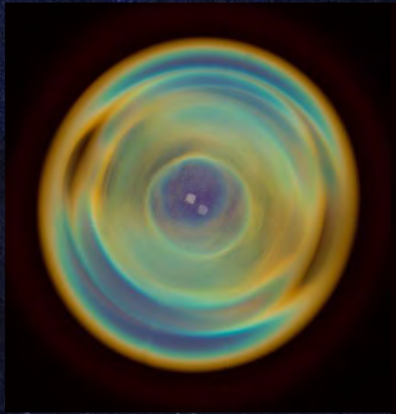

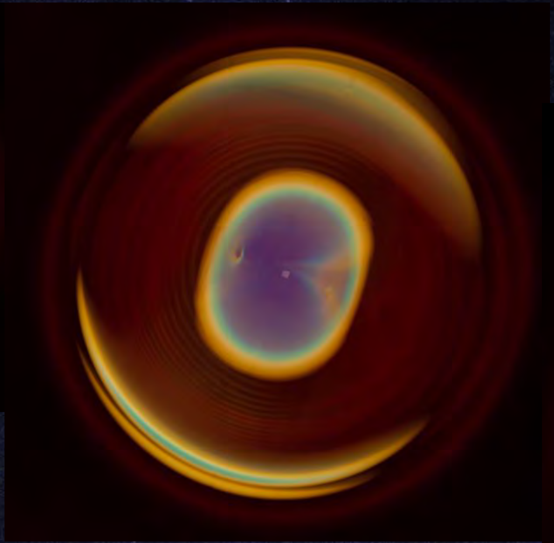

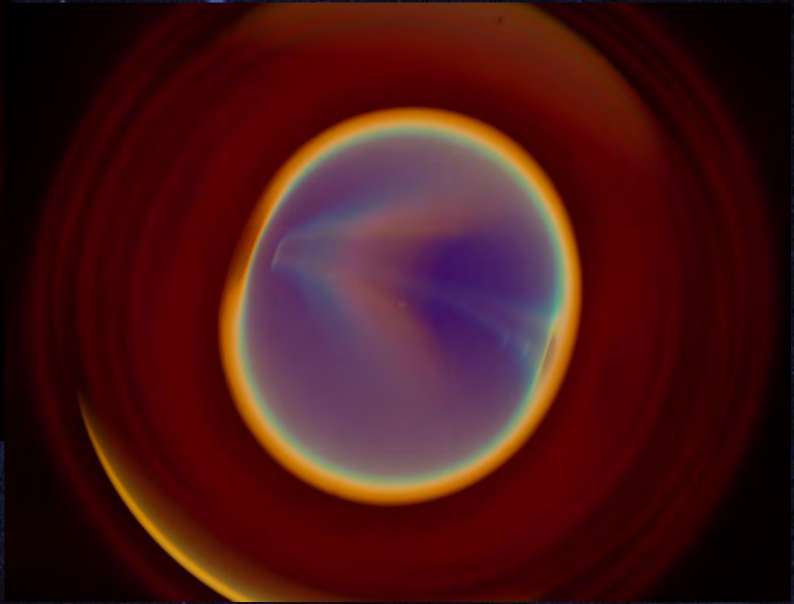

Image Credits:

Greg Abram – TACC

Francesca Samsel – CAT

Carson Brownlee - Intel

Markus Kunesch, Juha Jäykkä, Pau Figueras, Paul Shellard

Center for Theoretical Cosmology, University of Cambridge

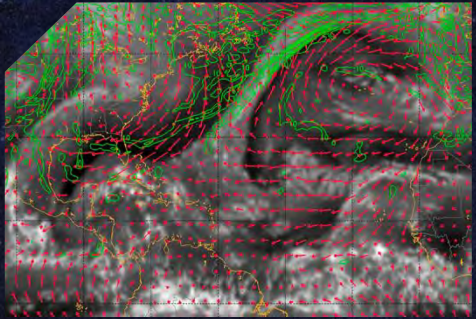

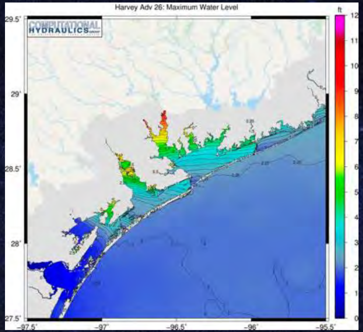

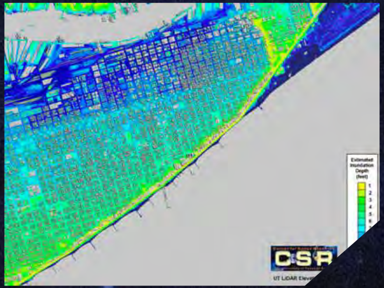

HARVEY

- Next Generation Storm Forecasting (with Penn State)
- Storm Surge Modeling (with Clint Dawson UT Austin)
- Preliminary river flooding and inundation maps (David Maidment UT Austin)
- Remote Image Integration and Assimilation (Center for Space Research, UT Austin)

MASSIVE DATA SET WORTHY OF ROSS ICE SHELF ITSELF

TACC partners with Lamont-Doherty Earth Observatory (LDEO) to host for one of the country's largest earth sciences data collections

- Managing hundreds of TB using Stampede2, Corral, and Ranch: storage, provenance, visualization, and public access
- Achieved 10x workflow speedup by moving to TACC (from 50 hrs down to 5 hrs for transfer and analysis tasks)

"...partnership...with TACC shows [it's] possible to manage...this level of data in a cost-effective, user-friendly and easily accessible manner..."

Image courtesy Oceanwide Expeditions.

TACC Press Release

PI Lingling Dong, Columbia University

XSEDE support to multidisciplinary, multi-institutional Rosetta project

"Using commodity HPC servers…the time to data-driven discovery is reduced and overall efficiency can be significantly increased." (Niall Gaffney, TACC)

RECORD ACHIEVED ON AI BENCHMARK

Graphic credit Andrej Karpathy

TACC, Berkeley, Cal Davis collaborate on large-scale AI runs

- Research demonstrating the potential of commodity hardware for AI
- Skylake ImageNet benchmark: (100 epochs, 11 min, 1024 nodes) -- fastest result at time of publication
- Knights Landing ImageNet benchmark (90 epochs, 20 min, 2048 nodes) – 3x faster than Facebook, with higher large-batch accuracy

Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, Kurt Keutzer

TACC Press Release

BRAIN TUMOR SEGMENTATION

- A team of researchers led by George Biros from The University of Texas at Austin scored in the top 25% of participants in the Multimodal Brain Tumor Segmentation Challenge 2017 (BRaTS'17) enabled by Stampede2 and other TACC resources.
- In the challenge, research groups presented methods and results of computer-aided identification and classification of brain tumors, as well as different types of cancerous regions.
- The team's method combined biophysical models of tumor growth with machine learning algorithms for the analysis of Magnetic Resonance imaging data of glioma patients.

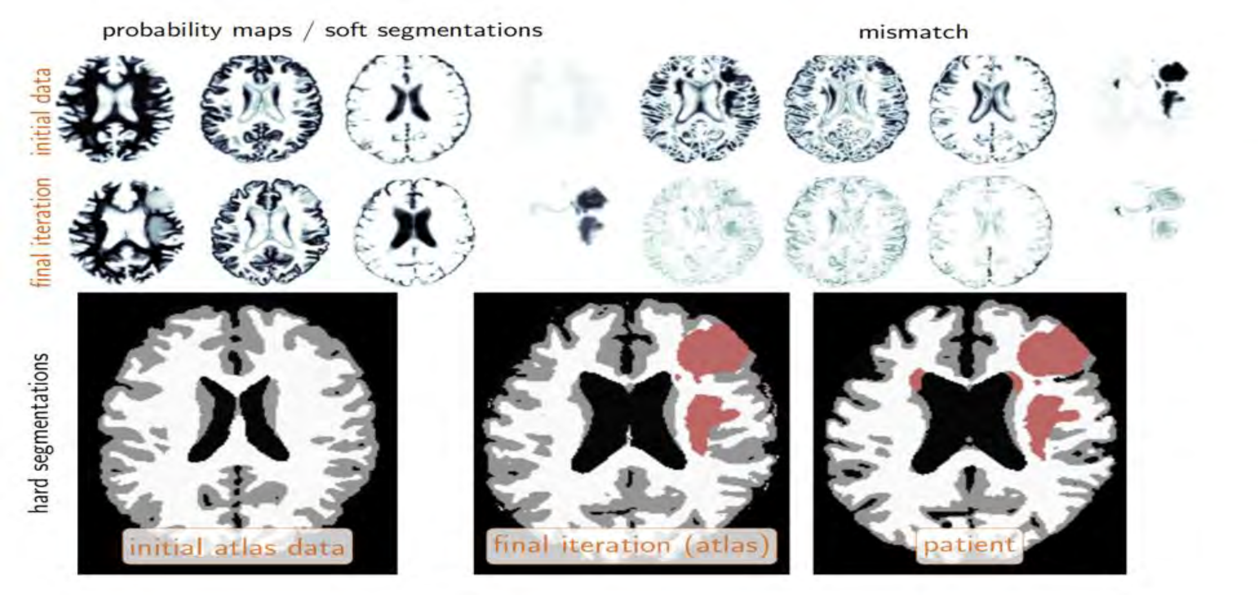

FRONTERA SYSTEM --- PROJECT

- *A new, NSF supported project to do 3 things:*
- Deploy a system in 2019 for the largest problems scientists and engineers currently face.
- Support and operate this system for 5 years.
- Plan a potential phase 2 system, with 10x the capabilities, for the future challenges scientists will face.

FRONTERA

FRONTERA SYSTEM --- HARDWARE

- Primary compute system: DellEMC and Intel
  - 35-40 PetaFlops Peak Performance
- Interconnect: Mellanox HDR and HDR-100 links.
  - Fat Tree topology, 200Gb/s links between switches.
- Storage: DataDirect Networks
  - 50+ PB disk, 3PB of Flash, 1.5TB/sec peak I/O rate.
- Single Precision Compute Subsystem: Nvidia
- Front end for data movers, workflow, API

FRONTERA

DESIGN DECISIONS - PROCESSOR

- The architecture is in many ways “boring” if you are an HPC journalist, architect, or general junkie.
  - We have found that the way users refer to this kind of configuration is “useful”.
- No one has to recode for higher clock rate. We have abandoned the normal “HPC SKUS” of Xeon, in favor of the Platinum top bin parts – the ones that are 205W per socket.
  - Which, coincidentally, means the clock rate is higher on every core, whether you can scale in parallel or not.
  - Users tend to consider power efficiency “our problem”.
  - This also means there is *no* air cooled way to run these chips.
- Versus Stampede2, we are pushing up clock rate, core count, and main memory speed.
  - This is as close to “free” performance as we can give you.

DESIGN DECISIONS - FILESYSTEM

- Scalable Filesystems are always the weakest part of the system.
  - Almost the only part of the system where bad behavior by one user can affect the performance of a *different* user.
- Filesystems are built for the aggregate user demand – rarely does one user stress *all* the dimensions of filesystems (Bandwidth, Capacity, IOPS, etc.)
- We will divide the "scratch" filesystem into 4 pieces
  - One with very high bandwidth
  - 3 at about the same scale as Stampede, and divide the users.
- Much more aggregate capability – but no need to push scaling past ranges at which we have already been successful.
  - Expect higher reliability from perspective of individual users
  - Everything POSIX, no "exotic" things from user perspective.

ORIGINAL SYSTEM OVERVIEW

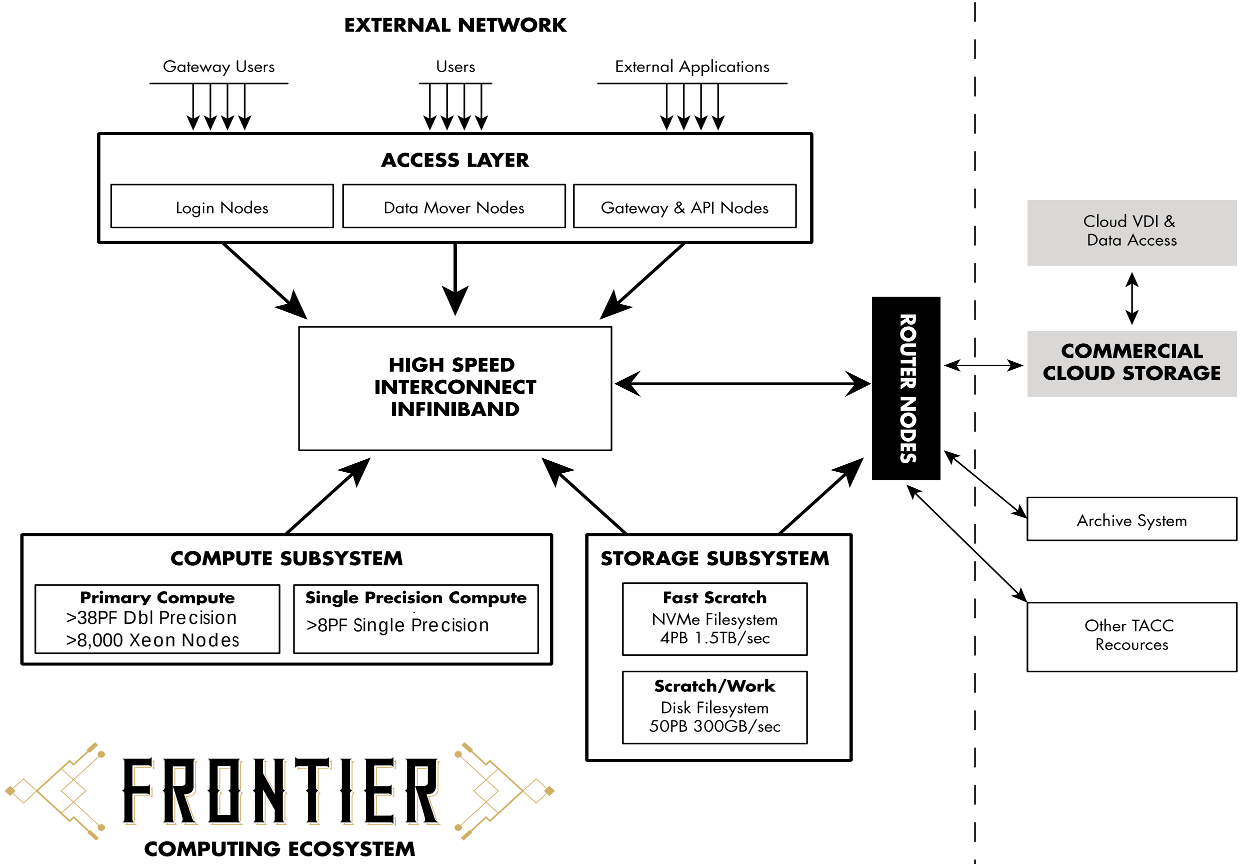

FRONTERA SYSTEM --- INFRASTRUCTURE

- Frontera will consume almost 6 Megawatts of Power at Peak
- Direct water cooling of primary compute racks (CoolIT/DellEMC)
- Oil immersion Cooling (GRC)
- Solar, Wind inputs.

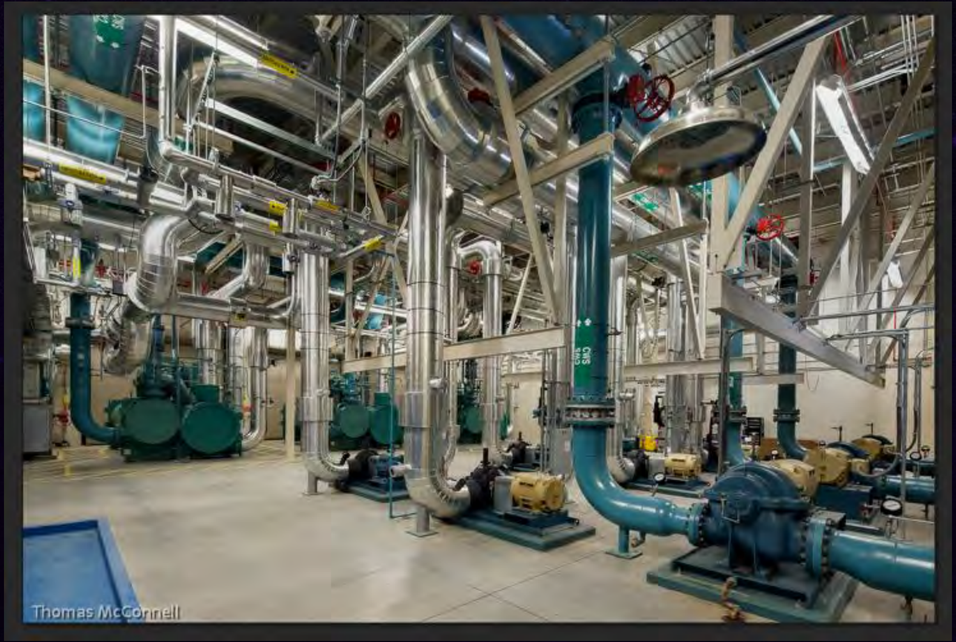

TACC Machine Room Chilled Water Plant

THE TEAM - INSTITUTIONS

- Operations: TACC, Ohio State University (MPI/Network support), Cornell (Online Training), Texas A&M (Campus Bridging)
- Science and Technology Drivers and Phase 2 Planning: Cal Tech, University of Chicago, Cornell, UC-Davis, Georgia Tech, Princeton, Stanford, Utah
- Vendors: DellEMC, Intel, Mellanox, DataDirect Networks, GRC, CoolIT, Amazon, Microsoft, Google

PHASE 2 PROTOTYPES

- Allocations will include access to testbed systems with future/alternative architectures
  - Some at TACC, e.g. FPGA systems, Optane NVDIMM, {as yet unnamed 2021, 2023}.
  - Some with partners – a Quantum Simulator at Stanford.
  - Some with the commercial cloud – Tensor Processors, etc.
- Fifty nodes with Intel Optane technology will be deployed next year in conjunction with the production system
  - Checkpoint file system? Local checkpoints to tolerate soft failures? Replace large memory nodes? Revive "out of core" computing? In-memory databases?
- Any resulting phase 2 system is going to be the result, at least in part, of actual users measured on actual systems, including at looking at, what they might actually *want* to run on.
- Eval around the world – keep close tabs on what is happening elsewhere (sometimes by formal partnership or exchange – ANL, ORNL, China, Europe).

STRATEGIC PARTNERSHIP WITH COMMERCIAL CLOUDS

- Cloud/HPC is *not* an either/or. (And in many ways, we are just a specialized cloud).
- Utilize cloud strengths:
  - Options for publishing/sustaining data and data services
  - Access to unique services in automated workflow; VDI (i.e. image tagging, NLP, who knows what. . .)
  - Limited access to *every* new node technology for evaluation
    - FPGA, Tensor, Quantum, Neuromorphic, GPU, etc.
  - We will explore some bursting tech for more "throughput" style jobs – but I think the first 3 bullets are much more important. . .

THE BROADER TACC ECOSYSTEM
DISCOVERY SCIENCE AT ALL SCALES

FRONTERA

Leadership/Discovery Science

Longhorn
IBM Power 9 +GPU
400+ Nvidia V100s
AI/ML/DL @ Scale

Testbeds
Catapult (Upgrade)
Non-Volatile Memory
Quantum
Future . . .

Existing TACC Computing Systems

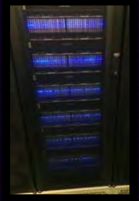

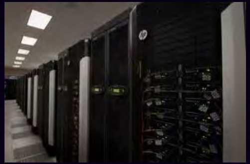

Existing TACC Storage Systems

DESIGNSAFE-CI

A NATURAL HAZARDS ENGINEERING COMMUNITY

Using Cyberinfrastructure to Reduce the Impacts of Natural Hazards

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY | TEXAS The University of Texas at Austin | TACC | RICE | Florida Tech

Disaster Resilience

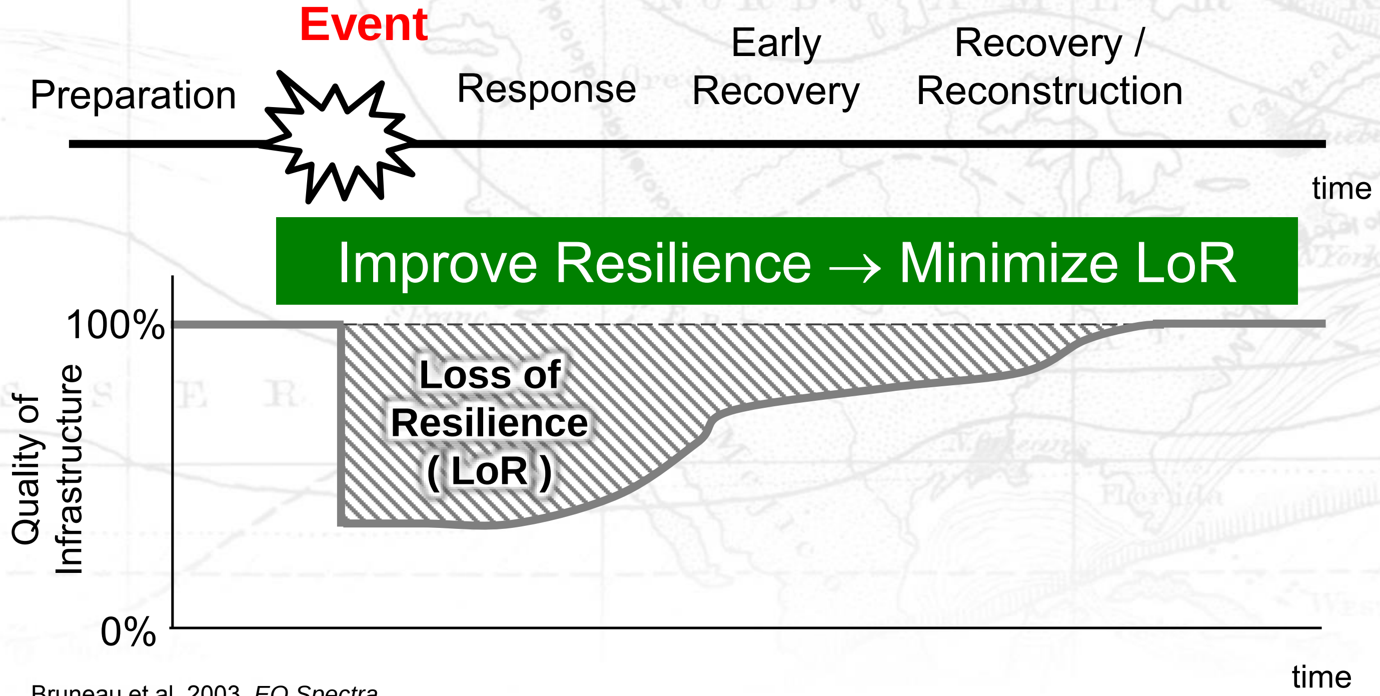

Bruneau et al. 2003, *EQ Spectra*

Improving Resilience

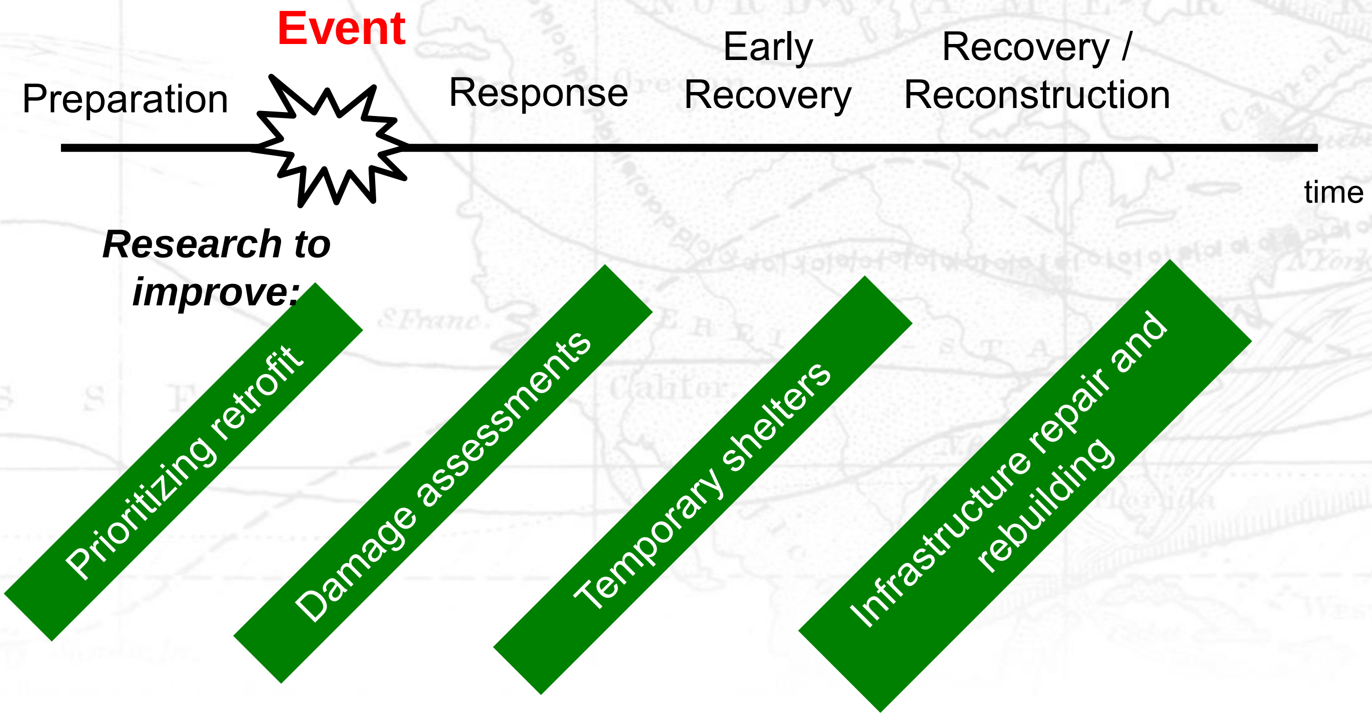

Improving Resilience - Impediments

- Limited availability of data
  - Damaging effects of previous disasters and the resulting response/recovery/rebuilding
  - Experimental data investigating engineering solutions
  - Simulation data predicting effects of natural hazards
- Why is little data available?
  - Researchers may collect data but do not share it broadly
  - After data are used by a researcher, it is “filed away”
  - No mechanism to reward data sharing

Improving Resilience - Impediments

- Integrating diverse datasets is cumbersome
  - Data sets from different fields are found in different places
  - Large data sets are hard to move
  - Visualization is often needed

DESIGNSAFE-CI
A NATURAL HAZARDS ENGINEERING COMMUNITY

What is DesignSafe?

- A web-based research platform that provides computational tools to manage, analyze, and understand critical data for natural hazards research

DesignSafe Vision

- A CI that is an integral part of research discovery
  - Support end-to-end research workflows and the full research lifecycle, including data sharing/publishing
  - Cloud-based tools that support the analysis, visualization, and integration of diverse data types
- Amplify and link the capabilities of natural hazards researchers in the US and abroad

NHERI: Natural Hazards Engineering Research Infrastructure

- NSF-funded, shared-use research infrastructure for research in natural hazards engineering
  - Network Coordinating Office (NCO)
  - Cyberinfrastructure (CI)
  - Seven experimental facilities (EF)
  - Post-disaster, rapid response research facility (RAPID)
  - Computational Modeling and Simulation Center (SimCenter)
- Replaces similar program for earthquake engineering (NEES) but expanded to include windstorms and associated hazards

NHERI Facilities

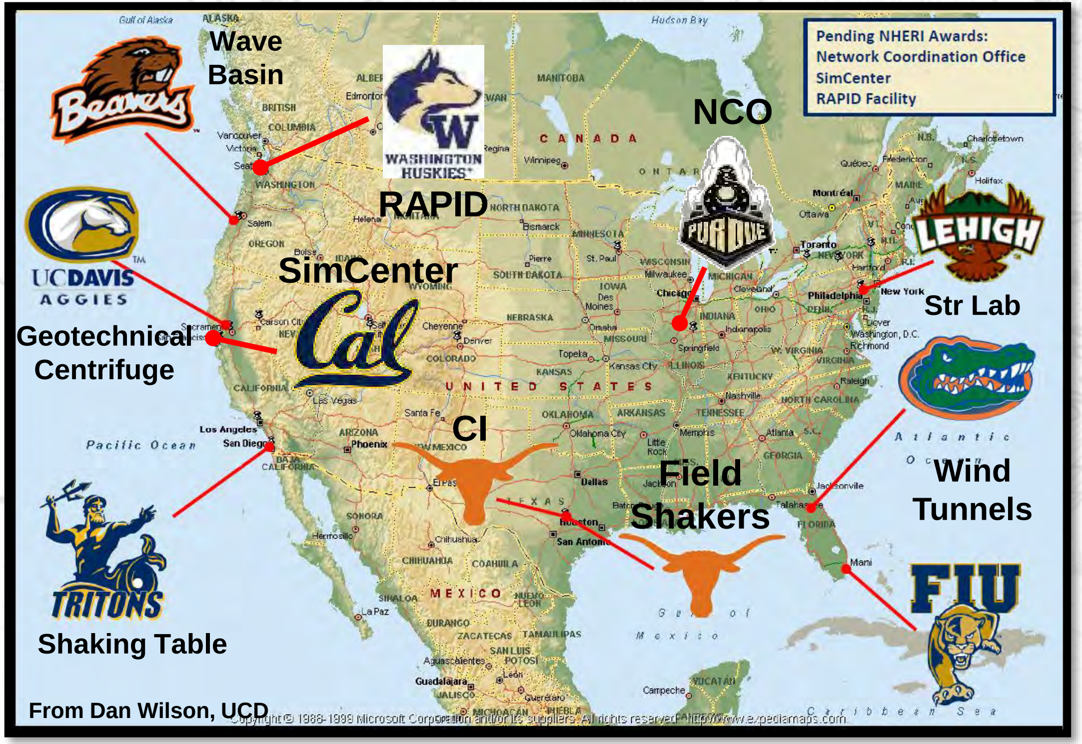

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY | TEXAS The University of Texas at Austin | TACC | RICE | Florida Tech

DesignSafe Components

- Research Workbench
  - Data Depot
  - Reconnaissance Portal
  - Discovery Workspace
- Learning Center
  - Training resources and student engagement
- NHERI Facilities
  - Access to information about all NHERI facilities
- NHERI Community
  - News and online Slack community

Data Depot Features

- Different areas:
  - My Data (Private)
  - My Projects (Semi-Private, Collaborative)
  - Published (Publicly accessible, curated)
  - Community Data (Publicly accessible, uncurated)
- Upload files/folders via computer, cloud service providers, or bulk transfer (Globus)
- Manage, preview files within Data Depot
- Data curation and publishing

My Projects

DESIGNSAFE-CI

NHERI: A NATURAL HAZARDS ENGINEERING RESEARCH INFRASTRUCTURE

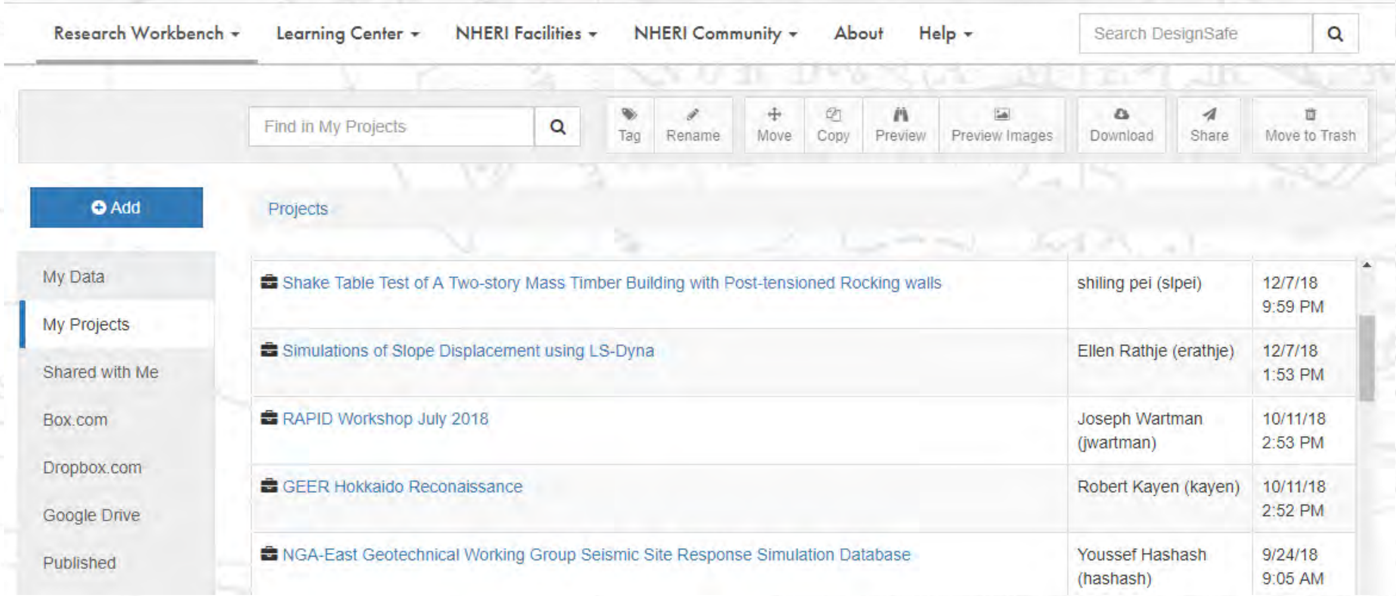

A space to share files/data/results with collaborators and to eventually publish for public use

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY

TEXAS The University of Texas at Austin

TACC

RICE

Florida Tech

Data Curation Philosophy

- Allow users to easily store, share, document, and publish data throughout the life of a research project
- Flexible data models and interactive curation
  - Allows researchers to decide how to represent their research
  - Support different types of data (experimental, simulation, field)
  - Consider what is needed for data to be understandable by others for data reuse

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY · TEXAS The University of Texas at Austin · TACC · RICE · Florida Tech

Published Projects

Add

- Published
- Community Data

Curation Tutorials
Curation Guidelines

Published PRJ-1693

PRJ-1693: CENTRIFUGE MODELLING OF VARIABLE RATE CONE PENETRATION IN LOW-PLASTICITY SILTS

PI	Boulanger, Ross — View Team Members	**DOI**	doi:10.17603/DS2J67J — Citation
Date of Publication	Apr/17/2018	**Award**	CMMI-1300518
Project Type	Experimental	**Keywords**	centrifuge, intermediate soils, silt, CPT
Associated Projects	Direct simple shear testing for silica silt and kaolin mixtures - PRJ-1689		

Description

This experimental program supports a broader project toward validating methods for the mechanics-based development of relationships between cone penetration test measurements and the engineering properties of intermediate soils.

Experiment ABP03s - Centrifuge Modelling of Variable Rate Cone Penetration in 80S20K Silt-Clay Mixture (80g)

ABP03s - Centrifuge Modelling of Variable Rate Cone Penetration in 80S20K Silt-Clay Mixture (80g)

Authors Price, Adam; Boulanger, Ross; DeJong, Jason;

DOI doi:10.17603/DS2DD46 Citation

Experimental Facility Center for Geotechnical Modeling, UC Davis

Experiment Type Centrifuge

Equipment Type 1m Radius Dynamic Geotechnical Centrifuge

Date of Publication Apr/17/2018

Description

Centrifuge test with in-flight variable rate cone penetration soundings in a slurry deposited silt-clay mixture of 80% silica silt and 20% kaolin clay by dry mass (80S20K) with a plasticity index of 6. The test was performed at a centrifugal acceleration of 80g. Cone soundings were performed using a 6-mm diameter cone penetrometer.

Accelerating Research: Data Re-Use

- ***Formal publishing*** of well-documented/valuable data sets for re-use must be recognized as scholarly work
- Data needs a permanent, ***digital location (DOI)*** similar to journal article, not just a URL
- Formally cite data ***in your reference list*** of your papers using DOI, citation language as indicated in DesignSafe

Experiment ABP03s - Centrifuge Modelling of Variable Rate Cone Penetration in 80S20K Silt-Clay Mixture (80g)

ABP03s - Centrifuge Modelling of Variable Rate Cone Penetration in 80S20K Silt-Clay Mixture (80g)

Authors Price, Adam; Boulanger, Ross; DeJong, Jason; **DOI** doi:10.17603/DS2DD46 Citation

Experimental Facility Center for Geotechnical Modeling, UC Davis
Experiment Type Centrifuge
Equipment Type 1m Radius Dynamic Geotechnical Centrifuge
Date of Publication Apr/17/2018

MAKE DATA COUNT

Citation

Price, Adam; Boulanger, Ross; DeJong, Jason, (2017), "ABP03s - Centrifuge Modelling of Variable Rate Cone Penetration in 80S20K Silt-Clay Mixture (80g)" , DesignSafe-CI [publisher], Dataset, doi:10.17603/DS2DD46

Reconnaissance Portal

Identifying Archived Datasets from Recon Events

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY | TEXAS The University of Texas at Austin | TACC | RICE | Florida Tech

Reconnaissance Portal

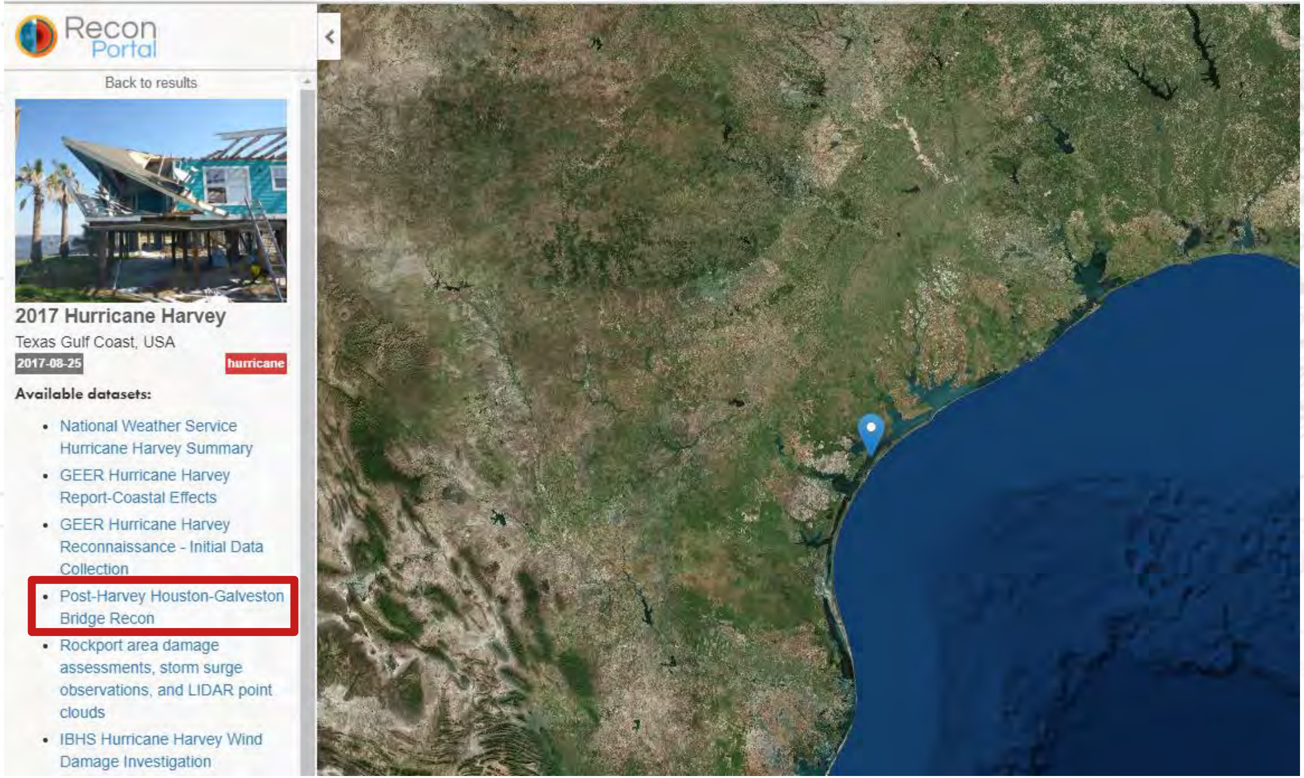

Recon Portal → Data Depot

Add

- My Data
- My Projects
- Shared with Me
- Box.com
- Dropbox.com
- Google Drive
- Published
- Community Data

Learn About 'My Projects'

Published / PRJ-1900

PRJ-1900: GEER HURRICANE IRMA - CAPE CORAL TO KEY WEST (INITIAL DATA COLLECTION)

PI	Stark, Nina (View Team Members)	**DOI**	doi:10.17603/DS2239D (Citation)
Date of Publication	Apr/19/2018	**Award**	NSF CMMI-1266418 (GEER)
Project Type	Other	**Keywords**	reconnaissance, geotechnical, coastal, Hurricane Irma 2017

Description

Hurricane Irma was a category 5 hurricane on the Saffir-Simpson hurricane wind scale. Irma developed from a tropical wave around the Cape Verde Islands. The National Hurricane Center started monitoring it on August 26, and it was classified as a tropical storm named Irma on August 30. Moving across the Atlantic Ocean, Irma increased in strength. On September 5, Irma was classified as a category 5 hurricane with wind speeds up to 175 mph (280 km/h). Irma made landfall in the U.S. on Cudjoe Key (near Big Pine and Summerland Keys) in the morning of September 10, still being a category 4 hurricane, and made a second landfall on Marco Island, south of Naples, on the same day as a category 3 hurricane. In preparation for Hurricane Irma, more than 6.5 million people were ordered to evacuate (http://www.pbs.org), 134 fatalities were associated to the storm, and damages were recorded of more than $50 billion (http://www.bbc.com/news/business-41231323; https://en.wikipedia.org/wiki/Hurricane_Irma). Two teams from the Geotechnical Extreme Events Reconnaissance (GEER) Association, supported by the National Science Foundation, were deployed to investigate geotechnical impacts of flooding, storm surge and wave forcing in Florida in response to Hurricane Irma in September of 2017. The teams worked collaboratively with federal, state, and local organizations in Florida. This initial data collection presents the field observations of the GEER team made during the field reconnaissance from September 24 to 28, 2017. The survey region extended along the coastal zone from Cape Coral to Key West.

Name	Size	Last modified
floridairma-coord.kml	378.1 kB	4/19/18 3:50 PM
GEER_report_#56.pdf	22.8 MB	4/19/18 3:48 PM
Photos	--	4/19/18 12:39 PM
READ ME_GEERIrma.pdf	305.8 kB	4/19/18 3:50 PM

Discovery Workspace

- Cloud-based tools for use in research
- Access to files in the Data Depot

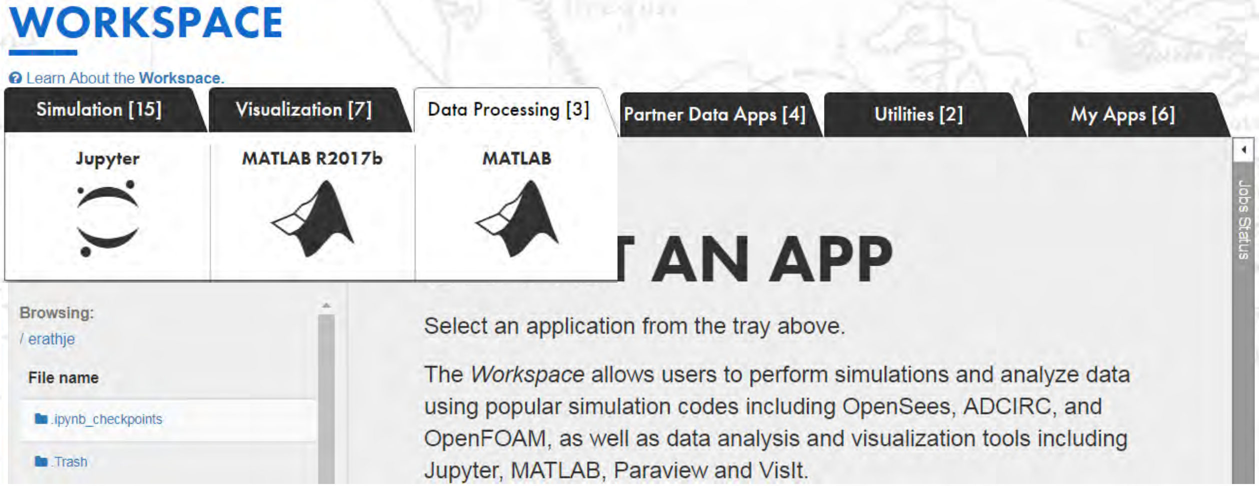

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY | TEXAS The University of Texas at Austin | TACC | RICE | Florida Tech

DesignSafe Discovery Workspace

- Data analysis in the cloud
  - Matlab: data analysis and plots
  - Jupyter : electronic notebook that supports Python and R
- Computational simulation codes
  - OpenSees: finite element code for structures and soil
  - ADCIRC: storm surge modeling
  - OpenFOAM: computational fluid dynamics
  - LS-DYNA: available via Bring Your Own License
- Visualization in the cloud
  - Potree: View and analyze point cloud data
  - QGIS: geospatial data analysis

Electronic Data Report

jupyter Data Report v 02 (read only)

Control Panel Logout

File Edit View Insert Cell Kernel Widgets Help

Python 3

Code CellToolbar

Cite

Data Report: Centrifuge Testing of a Circular and a Rectangular Embedded Structure During Base Shaking

Elnaz Esmaeilzadeh Seylabi [*1], Eva Agapaki [†2], Dimitris Pitilakis[‡3], Scott J. Brandenberg[§2], Jonathan P. Stewart[¶2], and Ertugrul Taciroglu[‖2]

[1]*Mechanical & Civil Engineering Department, California Institute of Technology, CA 91125*
[2]*Civil & Environmental Engineering Department, University of California, Los Angeles, CA 90095*
[3]*Department of Civil Engineering, Aristotle University of Thessaloniki, 54124, Thessaloniki, Greece*

[*]*Postdoctoral Researcher, elnaz@caltech.edu*
[†]*agapakieva@gmail.com*
[‡]*Assistant Professor, dpitilak@civil.auth.gr*
[§]*Associate Professor, sjbrandenberg@ucla.edu*
[¶]*Professor, jstewart@seas.ucla.edu*
[‖]*Professor, etacir@ucla.edu (Principal Investigator)*

1 Introduction

Seismic response of underground structures is a complex soil-structure interaction (SSI) problem in which two fundamental mechanisms are at play. Kinematic SSI is concerned with the motion of the structure in the presence of spatially variable ground motions and the interface pressures that develop as a result of different structural and free-field motions. Inertial SSI captures the soil reactions that develop to resist inertial forces associated with the acceleration of the structure. The kinematic component is generally considered to be more significant for buried structures, due to their modest mass and their confinement with the surrounding soil.

Published Analysis Scripts

jupyter 3b_app_flex_rev3_updated Last Checkpoint: 05/19/2017 (autosaved)

Control Panel Logout

File Edit View Insert Cell Kernel Widgets Help Python 3

Markdown CellToolbar

Application of Probabilistic Framework for Flexible Sliding Displacements

Vector (PGA, PGV) Approximation

Site Location: W -121.99 and N 37.18, Santa Clara County, California

Deaggregation Source: (https://earthquake.usgs.gov/hazards/interactive/) - Vs30 = 760 m/s

Mean and standard deviations of ground motions (ie. PGA and PGV) Source: NGA_Models_Version2.4.xls

MRE for PGA obtained from deaggragation Correlation coefficient between PGA and PGV, rhopgapgv = 0.6 Correlation coefficient between kmaxz and kvmaxz, rho_kmaxkvmax = 0.6 Logic tree consists of 27 branches: 9 for Ts and ky (correlated) and 3 for Tm

Reference: Rathje, E.M., Wang, Y., Stafford, P.J., Antonakos, G. and Saygili, G., 2014. Probabilistic assessment of the seismic performance of earth slopes. Bulletin of Earthquake Engineering, 12(3), pp.1071-1090.

Site Location

Out[4]:

DESIGNSAFE-CI A NATURAL HAZARDS ENGINEERING COMMUNITY | TEXAS The University of Texas at Austin | TACC | RICE | Florida Tech

DesignSafe Workflow Example

What addresses will be inundated on Galveston Island by storm surge from the impending hurricane?

- Compute storm surge water levels with ADCIRC
- Convert output to shapefile format
- Import results into a GIS along with elevation and property data
- Identify addresses that are inundated by simulated water levels

Kalpana python script used to convert ADCIRC output files to shapefiles. Executed within a Jupyter notebook in DesignSafe

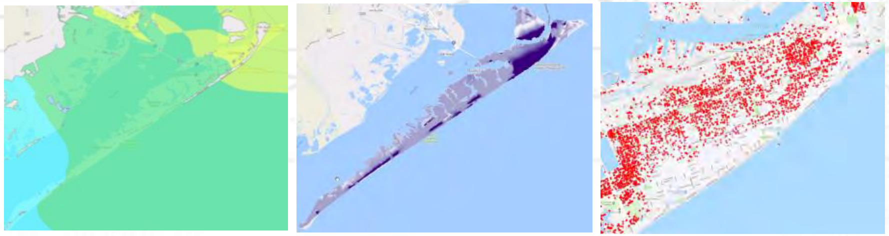

DesignSafe: Open for Business

www.designsafe-ci.org

- Capabilities available to the global natural hazards research community—account registration is free
- Working to expand our reach
  - Social science/urban planning/health care
  - Earth science

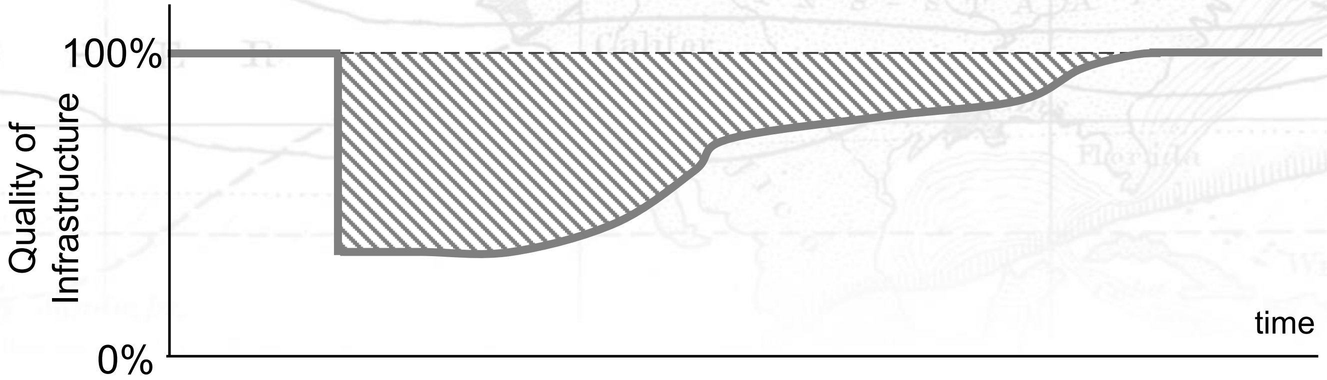

THANKS!!

- The National Science Foundation
- The University of Texas
- Peter and Edith O'Donnell
- Dell, Intel, and our many vendor partners
- Cal Tech, Chicago, Cornell, Georgia Tech, Ohio State, Princeton, Texas A&M, Stanford, UC-Davis, Utah
- Our Users – the thousands of scientists who use TACC to make the world better.
- All the people of TACC

- Humphry Davy, Inventor of Electrochemistry, 1812

- (Pretty sure he was talking about our machine).

Nothing tends so much to the advancement of knowledge as the application of a new instrument. The native intellectual powers of men in different times are not so much the causes of the different success of their labours, as the peculiar nature of the means and artificial resources in their possession.

Humphry Davy

PICTUREQUOTES.com


FRONTERA
TACC
NSF
TEXAS