July 10, 2018

XSEDE Overview

XSEDE

Extreme Science and Engineering Discovery Environment

John Towns
XSEDE Principal Investigator
jtowns@ncsa.illinois.edu

Motivation for XSEDE:

- Scientific advancement across multiple disciplines requires a *variety of resources and services*
- XSEDE is about increased productivity of the community and providing expanded capabilities
  - leads to more science
  - is sometimes the difference between a feasible project and an impractical one
  - lowers barriers to adoption
- XSEDE provides a *comprehensive eScience infrastructure* composed of expertly managed and evolving advanced *heterogeneous digital resources and services* integrated into a general-purpose infrastructure

XSEDE – accelerating scientific discovery

- XSEDE's Vision: *a world of digitally enabled scholars, researchers, and engineers participating in multidisciplinary collaborations while seamlessly accessing advanced computing resources and sharing data to tackle society's grand challenges.*

- XSEDE's Mission: *to enhance the productivity of a growing community of scholars, researchers, and engineers through access to advanced digital services that support open research by coordinating and adding value to the leading cyberinfrastructure resources funded by the NSF and other agencies.*

XSEDE Factoids: high order bits

- 5 year, US$110M project
  - pursuing additional funding via independent proposals
  - initial 5 year award: $121M project + ~$4.6M in supplements
    - plus $9M, 5 year Technology Investigation Service
      - separate award from NSF
- No funding for major hardware
  - coordination, support and creating a national/international cyberinfrastructure
  - coordinate allocations, support, training and documentation for >$100M of concurrent project awards from NSF
- ~90 FTE /~200 individuals funded across 19 partner institutions in 4 time zones
  - this requires solid partnering!

Vision/Mission: Enable Realizing Best Science

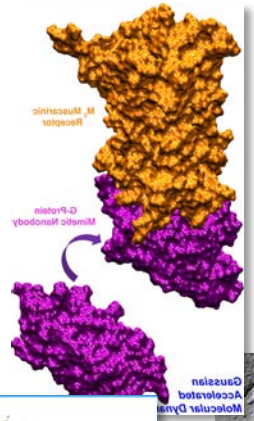

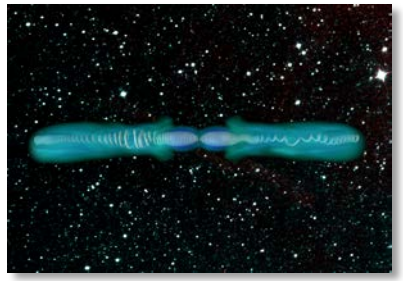

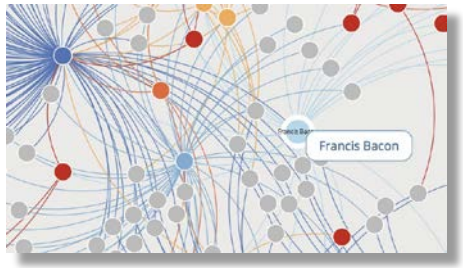

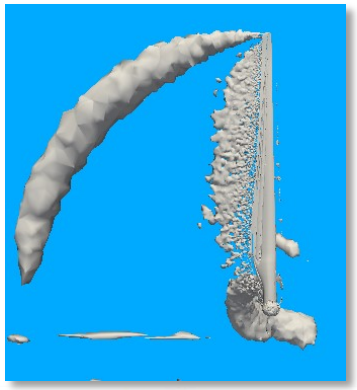

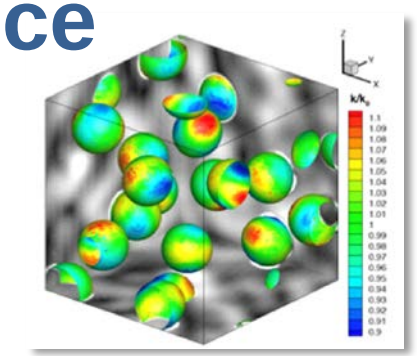

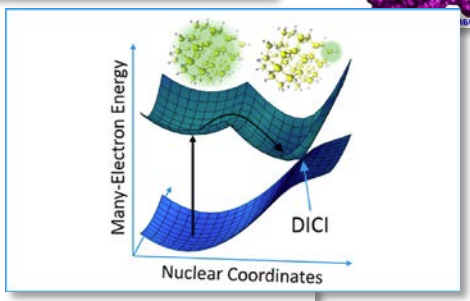

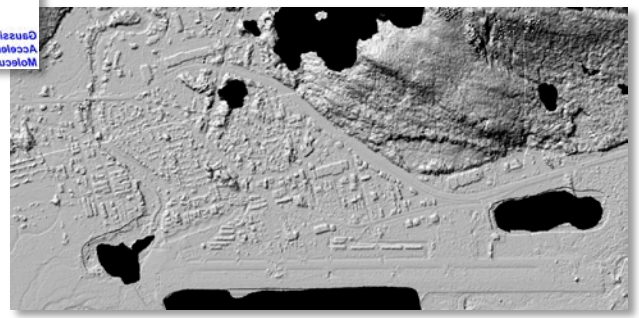

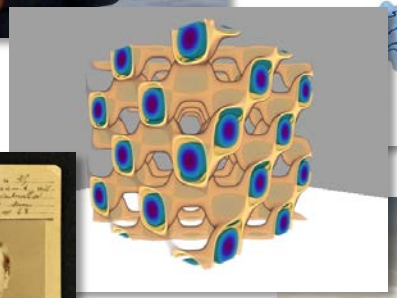

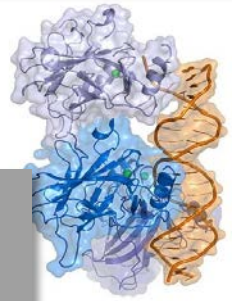

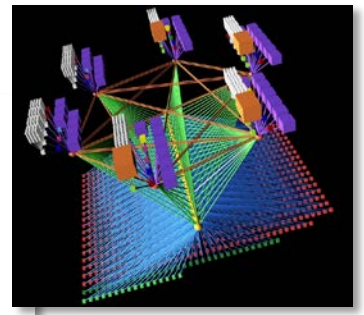

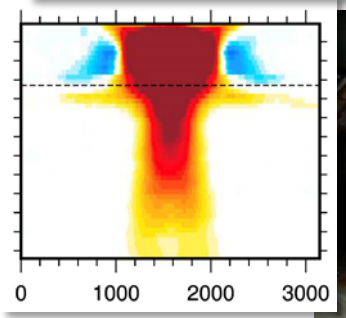

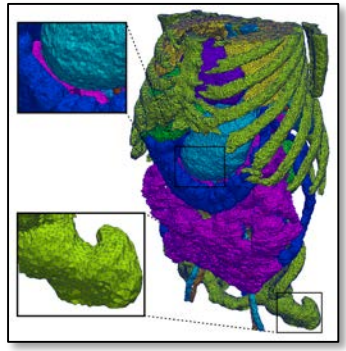

XSEDE

XSEDE's Distinguishing Characteristics:

World-class Leadership

Partnership led by

NCSA PSC TACC SDSC

Partners who strongly complement these CI centers with expertise in science, engineering, technology and education

SURA Southeastern Universities Research Association	The Shodor Education Foundation, Inc.	NICS	OU	Oklahoma State University
Ohio Supercomputer Center An OH-TECH Consortium Member	Purdue University	University of Arkansas	USC University of Southern California	NCAR
The University of Chicago	Indiana University	Cornell University	Georgia Tech	The University of Georgia

XSEDE

Total Research Funding Supported by XSEDE to Date

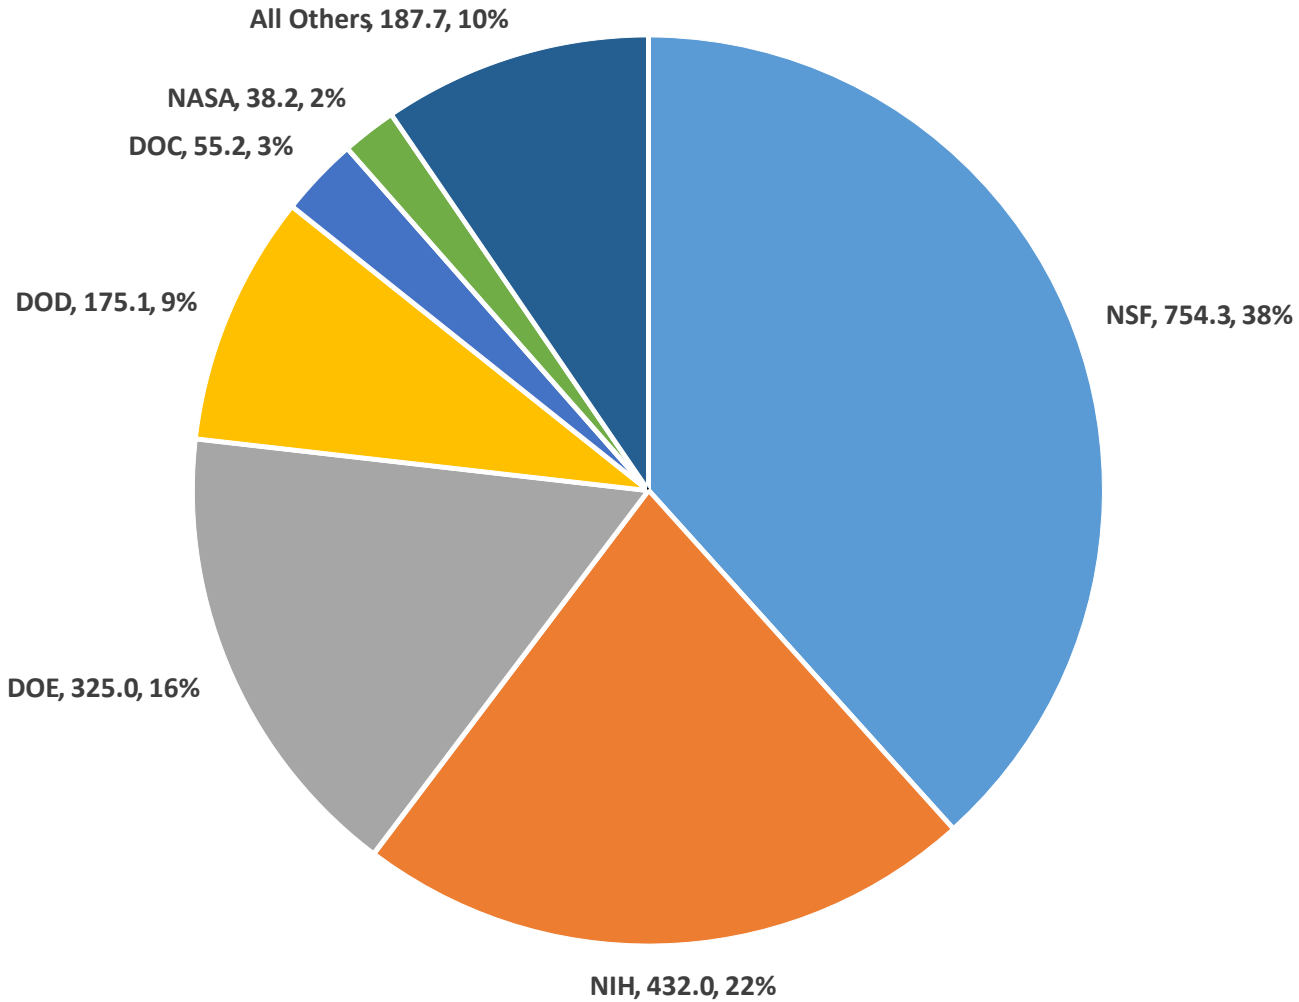

$1.97 billion in research supported by XSEDE 2.0

September 2016 - April 2018

Research funding only. XSEDE leverages and integrates additional infrastructure, some funded by NSF (e.g. "Track 2" systems) and some not (e.g. Internet2).

XSEDE offers efficient and effective integrated access to a variety of resources

- Leading-edge distributed memory systems
- Very large shared memory systems
- High throughput systems, including Open Science Grid (OSG)
- Support for VM's and containers and HPC Cloud
- Visualization engines
- Accelerators like GPUs and Xeon PHIs
- Extensive library of research applications

Many scientific problems have components that call for use of more than one platform.

XSEDE User Portal: THE User Site
portal.xsede.org

- XSEDE User Portal (XUP) is designed to be the only site a user needs to use XSEDE
- XUP presents information relevant to users
  - user info is easier to find
  - XUP also provides dynamic data about XSEDE systems
  - capabilities to manage usage, files, data
- As a user you can
  - request an allocation, and manage allocations
  - sign up for training
  - request help
  - manage files and data, and much more!
- Portal provides single sign-on to all XSEDE resources

Current XSEDE Compute Resources

- Stampede2 @ TACC
  - 12.8 PFLOPS (PF) Dell Cluster w/ Intel Knights Landing, Skylake
- Comet @ SDSC
  - 2.1 PF cluster w/GPUs
- Bridges @ PSC
  - 1 PF w/ large memory (3 TB and 12 TB) and GPU nodes
- XStream @ Stanford
  - 1.0 PF GPU Cray CS-Storm cluster
- SuperMIC @ LSU
  - 925 TF Dell Cluster w/ GPUs and Xeon PHIs
- Jetstream @ Indiana
  - 516 TF HPC Cloud
- Wrangler @ TACC
  - 62 TF data analytics system
- Open Science Grid
  - 160,000 CPU cores

https://www.xsede.org/web/xup/resource-monitor

Current XSEDE Data and Software Resources

- Storage
  - –Ranch @ TACC
    - •61 PB tape
  - –Pylon @ PSC
    - •10 PB disk
  - –Wrangler @ TACC
    - •10 PB disk
  - –Data Oasis @ SDSC
    - •4 PB tape
  - –Jetstream Storage
    - •1 PB disk

https://www.xsede.org/web/xup/resource-monitor

- Software: 100s of titles
  - –domain software
    - •chemistry, CFD, bioinformatics, physics, astronomy, biology, engineering, statistics,…
  - –tools
    - •middleware, visualization, scripting, performance analysis, data storage and management, …
  - –compilers and libraries
    - •most languages supported, math libraries, machine learning, …

https://portal.xsede.org/software#

ECSS is an Allocated Service

- Extended Collaborative Support Service
  - improves the productivity of the XSEDE user community through meaningful collaborations
- Expert staff can be requested for collaborations lasting months to a year
  - requests made through the XSEDE allocation system when requesting computational resources
- Typical collaborations require 20-25% staff time for one year
- Critical mass engenders success
  - 26 FTEs (64 individuals) for project work at 9 sites
    - +4 ECSS Affiliates (skilled volunteers)
  - advanced degrees in a variety of science and technology fields
    - some staff co-author publications or write proposals with PI team, a few are later funded by PI team

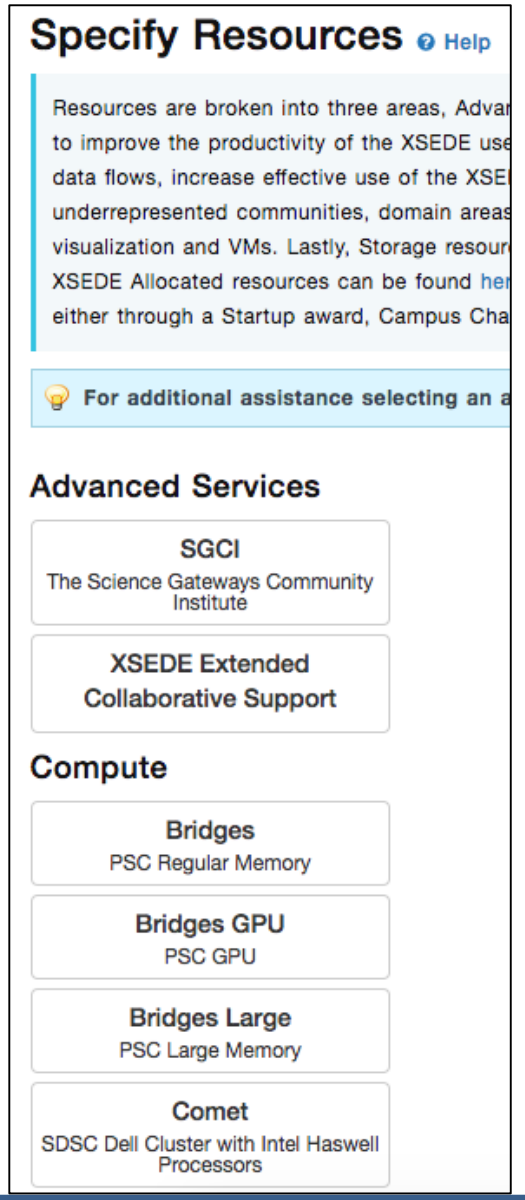

More information at: www.xsede.org