NSF XSEDE proposal: What? Why? Who?

Sohrab Ismail-Beigi
Becton 307

Departments:
Applied Physics
Physics
Mech. Eng. & Mat. Sci.

Outline

- Who am I?
- Why am I here?
- What are my computational usages & requirements?

- What is XSEDE?
- Types of XSEDE requests

- What does an XSEDE grant (request) look like?

- What is a typical outcome?

Outline

- **Who am I?**
- **Why am I here?**
- **What are my computational usages & requirements?**

- What is XSEDE?
- Types of XSEDE requests

- What does an XSEDE grant (request) look like?

- What is a typical outcome?

Who am I?

- Theoretical/computational physicist
- Study electronic behavior in solid materials
- At the level of individual atoms and electrons
- E.g. : electron distribution across
  an interface between metal
  oxides

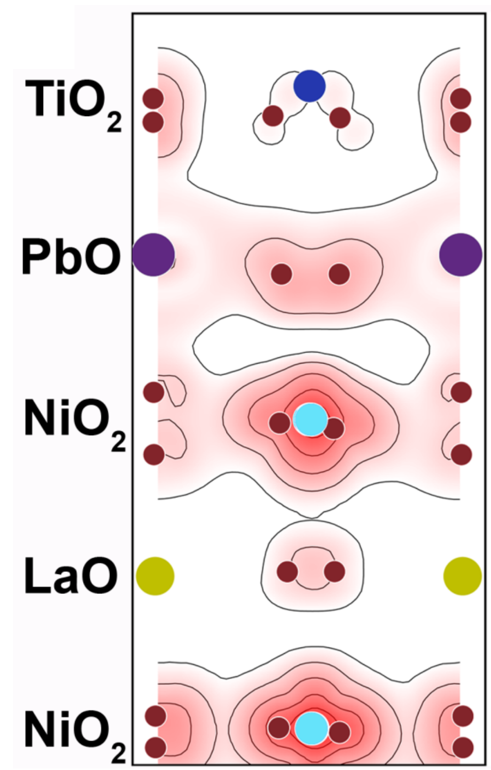

- Solve partial differential equations with large dimensions (quantum mechanics)

- Need high speed computation to crunch lots of numbers
  - Linear algebra, matrix diagonalization, Fourier transforms

Why am I here?

1. My group uses parallel high-speed computation for research

2. My group makes extensive use of Yale HPC clusters

3. We have been requesting and using
   XSEDE resources continuously since 2009

Our needs & modus operandi

- Primary parallel software is open source
- Fortran 90
- Use only 3 - 4 standard libraries:
  - FFTW, Blas, Lapack (ScaLapack)

- Typical production job uses 10 - 100 processors
- Each job is a single MPI parallel run using all processors

- Run **many** such calculations on a material and variants
  - Not few huge runs to get a few numbers
  - Many smaller runs to get detailed relationships
  - Variant calculations are human generated

- Burn through ~ 7-10 million CPU hours / year (whole group)

Our computer use breakdown

Group-wide

- ≈ 75% of compute time uses Yale HPC
- ≈ 25% of compute time

So, why do we use XSEDE?

1. Extra CPU time never hurts!
2. XSEDE supercomputers are large: queue times are short
   - Can get work done predictably and quickly
3. XSEDE supercomputers are at least as fast as Yale HPC
4. XSEDE support is quite good: dedicated team of HPC support staff who will usually work with you over many iterations to solve your issues
5. XSEDE supercomputers are quite reliable: almost no unannounced down time

Outline

- Who am I?
- Why am I here?
- What are my computational usages & requirements?

- **What is XSEDE?**
- **Types of XSEDE requests**

- What does an XSEDE grant (request) look like?

- What is a typical outcome?

What is XSEDE?

- Used to be TeraGrid
- Network of NSF-funded supercomputers
- Centered at NCSA (U. of Illinois Urbana-Champaign)
- Other members: Texas Austin, U. of Illinois, Ohio State, Pittsburgh & San Diego supercomputing, …

- Multiple types of resources
  - HPC parallel clusters (some with GPUs, memory, ...)
  - GPU clusters
  - Distributed high-throughput (embarrasingly parallel)
  - Visualization
  - Data storage
  - File management

Note: my group has only used the HPC resources, so that is the only thing I can competently talk about

What is XSEDE?

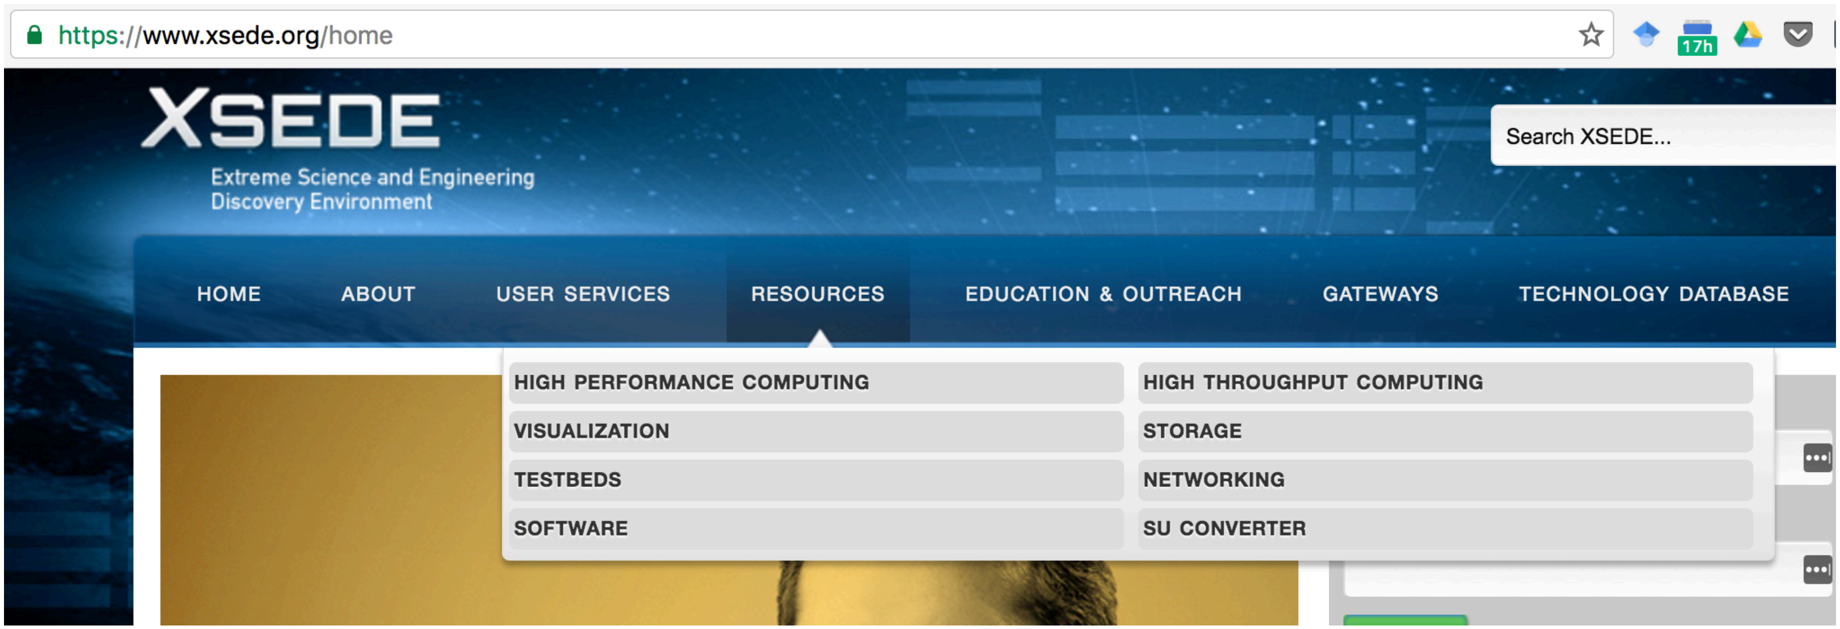

What is XSEDE?

https://www.xsede.org/resources/overview

Name	Site	Manufacturer / Platform	Machine Type	Peak Teraflops	Disk Size (TB)	Availability
Stampede User Guide	UT Austin	TACC Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors (Stampede)	Cluster	9600.0	14336.0	Production through 2017-09-30
Comet User Guide	SDSC	SDSC Dell Cluster with Intel Haswell Processors (Comet)	Cluster	2000.0	7000.0	Production through 2019-01-30
XStream User Guide	Stanford U	Stanford University GPU Cluster (XStream)	Cluster	1001.7	1400.0	Production
SuperMIC User Guide	LSU CCT	LSU Cluster (superMIC)	Cluster	925.0	840.0	Production
Bridges Regular Memory User Guide	PSC	PSC Regular Memory (Bridges)	Cluster	894.6		Production
Bridges Large Memory User Guide	PSC	PSC Large Memory Nodes (Bridges Large)	Cluster	894.6		Production
Jetstream User Guide	Indiana U	IU/TACC (Jetstream)	Cluster	516.1	1920.0	Production through 2019-11-30
Gordon Compute Cluster User Guide	SDSC	SDSC Appro with Intel Sandy Bridge Cluster (Gordon Compute Cluster)	Cluster	341.0	1628.0	Production through 2017-03-31
Wrangler User Guide	UT Austin	TACC Data Analytics System (Wrangler)	Cluster	62.0	5000.0	Production through 2019-01-30
Greenfield		PSC HP Superdome and				

What is XSEDE?

https://www.

Name

Stampede
User Gu

Comet
User Gu

XStream
User Gu

SuperMIC
User Gu

Bridges Re
Memory
User Gu

Bridges La
User Gu

Jetstream
User Gu

Gordon Co
Cluster
User Gu

UT Austin Stampede

Hostname	stampede.tacc.xsede.org
Site	tacc.xsede.org
Organization	University of Texas at Austin
Descriptive Name	TACC Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors (Stampede)
Manufacturer	Dell
Platform	Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors
CPU Type	Intel Xeon E5-2680
Machine Type	Cluster
Operating System	Linux (CentOS)
Contact	XSEDE Help Desk
Processor Cores	102400
Nodes	6400
Memory	200 TB
Peak Performance	9600 TFlops
Disk	14336
Primary Storage	

Wrangler User Guide	Austin	TACC Data Analytics System (Wrangler)	Cluster	62.0	5000.0	through 2019-01-30
Greenfield		PSC HP Superdome and				

Requesting XSEDE resources

- To get compute time on XSEDE, request an “allocation”
- Basically 4 types of allocations

Alloc type	CPU hours	Length	Wait time
Trial	~ 1K	6 months	~1 day
Startup	~ 10K – 50K	12 months	~ weeks
Education (for classes)	~ few K /student	12 months	~ weeks (?)
Research	~ 0.1M to some M *	12 months	Quarterly submission windows

** Asking for > 20M could bump up to a different category (e.g., NSF PRAC)*

First time XSEDEr

- Start with a Trial or Startup allocation
- Designed for testing/timing of XSEDE systems
  - Suitable for your needs?
  - Gather timings and parallel scaling data

- Then submit a Research proposal
  - Research proposals *must* have timing/scaling data
  - The data is taken seriously
  - Must get Startup before Research alloc. (official rule)

  - Can submit Research proposal while Startup is active

Outline

- Who am I?
- Why am I here?
- What are my computational usages & requirements?

- What is XSEDE?
- Types of XSEDE requests

- **What does an XSEDE grant (request) look like?**

- What is a typical outcome?

Education & Startup allocations: content

These are pretty short and sweet

1. Basic information about PI
2. Grants data supporting PI research
3. Abstract of research
4. Standard 2 page NSF CV
5. Resource request: which system, how much

Research allocations: content

Similar to a standard NSF research grant (but less painful ☺)

Key document	Page limit	Contents or main idea
Main text (NSF style: 1 inch margins, min. font size, …)	10 or 15 *	• Research justification • Research plan • Request & justification • Other compute support
Code performance & scaling	5	• Summary of code function • Scaling/timing data
CV	2 page NSF	
Progress report	3	*Renewals only*: results from prior XSEDE support
Publications resulting	none	*Renewals only*: publications from prior support

** Allocation request < 10M = 10 pages , > 10M request = 15 pages*

Main text: what's inside?

This is the document most reviewers look at most carefully

Things it must get across

1. Importance of research in context of state of art
2. Research plan: what questions to be answered and how
3. Resources needed to execute research plan
   and justification of the resources needed

In my group's requests, we literally have three separate sections dealing with these three issues

Main text: quick walk through PDF

(show & page through 2015 successful research request PDF

First 10 pages)

Main text example: importance

1 of our 7 projects: why do we do it?

Rare-earth nickelate oxides – Recently, piezoelectronic transistor (PET) devices have been proposed as an alternative to standard field-effect transistors [5]. PET functioning derives from the change of resistance of a material (used as the conducting channel) due to mechanical strain. Materials exhibiting strong dependence of resistance on applied strain (piezoresistive or PR materials) are therefore needed for efficient PET devices. Perovskite rare-earth nickelate oxides $R\text{NiO}_3$, where R is a rare-earth atom, may be suitable candidates for PR elements since they exhibit metal-insulator transitions with changes in resistivity of several orders of magnitude. The transition can be controlled by temperature, strain, and choice of R element. We wish to evaluate the suitability of several $R\text{NiO}_3$ thin film structures, such as SmNiO_3 and NdNiO_3, for application in PET architectures. To this end, we need to elucidate the precise dependence of crystal, electronic and magnetic structure of these materials on applied strain.

Main text example: what & how

1 of our 7 projects: questions & how to answer them

Rare-earth nickelate oxides – Our goal is to evaluate whether $NdNiO_3$, $SmNiO_3$, or their alloys $Nd_{1-x}Sm_xNiO_3$ are suitable piezoresistive materials for PET applications. First, we need to carefully benchmark the theoretical apparatus to make sure that we can correctly reproduce the experimentally observed bulk crystal structures and band gaps. For this purpose, we plan to perform DFT+U calculations using several exchange-correlation functionals to select the best choice. Hubbard U effects will be considered for both Ni $3d$ and O $2p$ states. After this necessary benchmarking, we will investigate the effect of various epitaxial strain modes on the electronic properties of $SmNiO_3$ and $NdNiO_3$ and several $Sm_xNd_{1-x}NiO_3$ structures. Specifically, we wish to know the dependence of the band gap on strain and find the epitaxial strain and composition combinations so that the electronic behavior (e.g., band gap) is most sensitive to strain.

Main text example: request & justification

1 of our 7 projects: what we need and why

Rare-earth nickelate oxides — To benchmark DFT+U calculations for bulk nickelate oxides, we will perform 80-atom cell calculations for bulk $NdNiO_3$ and $SmNiO_3$. For each system, we will consider LDA, PBE, and PBEsol exchange-correlation functionals. For each functional, we will use the DFT+U method with the Hubbard $U = 0$ eV and up to three positive U values. Since we examine Hubbard parameters on Ni $3d$ and O $2p$ states, we need to consider up to 16 combinations of U values (4 values for Ni and 4 values for O). Each relaxation requires on average 1,000 CPU hours. Therefore, for benchmarking stage we will need approximately 2×3×16×1,000=96,000 CPU hours. After identifying the best exchange-correlation approximation and U values, we will consider 5 mixed $Sm_xNd_{1-x}NiO_3$ structures with $x \in [0, 1]$. For each structure, we will perform 25 calculations with applied strain (5 calculations for in-plane strain times 5 for out-of-plane strain). This stage will require 5×5×5×1,000=125,000 CPU hours.

The total requested computational time for this project is 220,000 CPU hours.

Scaling and performance: quick walk through PDF

(show & page through 2015 successful research request PDF

Pages 11-14 of PDF)

Scaling and performance

- They are ***serious*** about having this document
- It is required
- And they look at the results

One needs to time and benchmark ones software
on their computers or near equivalents

It is not required that you show perfect performance

- You can still get resources
- They can recommend (or you can ask for) extended collaborative help from their experts to improve performance

Other documents: PDF walk through

- References : pages 15 - 17
- Progress report (prior support): pages 18 - 20
- Publications resulting (prior support): page 21
- CV: pages 22 - 23

Outline

- Who am I?
- Why am I here?
- What are my computational usages & requirements?

- What is XSEDE?
- Types of XSEDE requests

- What does an XSEDE grant (request) look like?

- **What is a typical outcome?**

Our most recent experience: fall 2015

SUBMISSION PERIOD	ALLOCATION BEGIN DATE
Dec 15 thru Jan 15	April 1
Mar 15 thru Apr 15	Jul 1
Jun 15 thru Jul 15	Oct 1
Sep 15 thru Oct 15	Jan 1

- Entire group started writing proposal pieces ~ Sept 15
- Assembly and proofreading done by ~ Oct 10
- Submitted research request Oct 15
- Received approval email decision Dec 17 2015 ☺

… but what did it say?

Our most recent experience: fall 2015

Renewal Approved

Opportunity Information

Resources

1. SDSC Appro with Intel Sandy Bridge Cluster (Gordon Compute Cluster)

Approved	1,613,053.00 SUs	Comments: (none)
Recommended	2,280,000.00 SUs	Comments: (none)
Requested	2,280,000.00 SUs	Comments: Having LAPACK, SCALAPACK, BLAS, FFTW, Fortran 90, C, C++, python, perl is very useful. We need queues that go longer than 24 hours. If PWSCF (quantum espresso) is already optimally compiled, then this is extremely useful to us.

2. SDSC Medium-term disk storage (Data Oasis)

Approved	500.00 GB	Comments: (none)
Recommended	500.00 GB	Comments: (none)
Requested	700.00 GB	Comments: (none)

- 30% cut
- Similar but slightly higher than prior years (~ 20% cuts)

Our most recent experience: fall 2015

- Why the cut?
- 2 reviewers were quite positive and didn't recommend cuts

- Problem: oversubscription
  i.e., too many researchers and not enough computers
  i.e., computers are never big enough

Meeting Comments

The total number of SUs requested was a factor of 3 times greater than available on all systems. After the usual merit-review criteria were applied by the assigned reviewers, and the panel-recommended allocations were determined, the totaled allocations for the cluster systems were found to be oversubscribed by 173M SUs. It was necessary to adjust the recommended allocations to fit within the budget of available SUs according to the formulation in section 6.4.1 of the XSEDE Allocations Policy document (www.xsede.org/web/guest/allocation-policy). Every panel-recommended allocation (derived from the review-recommended allocation during the panel discussion) was included in this reduction. Reductions of up to 50% were imposed on some recommended allocations to obtain the awarded allocation, depending upon the portion of NSF funding and the size of the recommended allocation.

Please note that over subscription, allocating above the available SUs, causes ...

Show More

Our most recent experience: fall 2015

- Our group is pretty happy overall with this

- It is still a good chunk of CPU time

- We are privileged: have Yale HPC (i.e., Yale $$$) to rely on

  How people deal at other schools? No idea…

Summary

- XSEDE is NSF's supercomputing infrastructure

- Requesting time is like writing a mini NSF grant
  - Easier than regular NSF grant but still a little painful

- It is free so don't complain too much

- You might get cut by ~ 20% – 30%

- Depending on your needs, it could be
  - Life saver
  - Workhorse
  - Padding
  - Extra candy

Questions?

(If you don’t ask, I can’t tell you that I don’t know ☺)