



XSEDE Resource Allocation System XTRAS

[XTRAS](#) / [Request](#)

Cryo-electron microscopy structure determination using the COSMIC2 science gateway

Submission for XRAC - June 2017 [View Opportunity](#)

Actions:

[View Reviews](#)

[Actions](#)

Number MCB170058

Status Approved

Type New

Abstract — Structural biology is in the midst of a revolution. Instrumentation and software improvements have allowed for the full realization of cryo-electron microscopy (cryo-EM) as a tool capable of determining atomic structures of protein and macromolecular samples. These advances open the door for solving new structures that were previously unattainable, which will soon make cryo-EM a ubiquitous tool for structural biology worldwide, serving both academic and commercial purposes. However, despite its power, new users to cryo-EM face significant obstacles. One major barrier consists of the handling of large datasets (10+ terabytes), where new cryo-EM users must learn how to interface with the Linux command line while also dealing with managing and submitting jobs to high performance computing resources. To address this barrier, we have developed the COSMIC2 Science Gateway as an easy, web-based, science gateway to simplify cryo-EM data analysis using a standardized workflow to run on XSEDE's [1] supercomputers. This gateway will lower the barrier to high performance computing tools and facilitate the growth of cryo-EM to become a routine tool for structural biology. This gateway is operational and will be open for beta testing shortly, requiring that we provide users with XSEDE SUs through our science gateway as a community

account. Therefore, we are requesting allocations on Comet, Comet GPU, XStream, and Oasis to provide up to 200 users access to cryo-EM software on supercomputer resources through the COSMIC2 science gateway.

Keywords — science gateway, structural biology, cryo-EM

Fields of Science

Biochemistry and Molecular Structure and Function primary

Personnel

PI

Michael Cianfrocco <mcianfro@umich.edu>

University of Michigan

315-264-1434

Supplement Incomplete

Actions:

Edit

Delete

Comment/Justification

We are requesting ECSS support to help benchmark RELION2 GPU performance on K80 and P100 GPUs.

Resources

1. XSEDE Extended Collaborative Support

Requested	1.00 [Yes = 1, No = 0]	Comments: (none)
What do you want to accomplish with the help of expert staff? Have you already done any work on this aspect of your software?		
Please enter your answer below: We are requesting ECSS to benchmark RELION2 GPU performance on K80 vs. P100 GPUs, and on single and multiple nodes on Comet GPU.		
How would the success of this collaboration benefit your project?		

Please enter your answer below: This benchmarking information will be critical for us in our COSMIC2 science gateway so that we can automatically distribute user jobs to GPUs in an effective manner

Which member(s) of your team would collaborate with ECSS staff?

Please enter your answer below: Michael Cianfrocco

Have you had significant interaction on previous projects related to your current proposal or discussed your extended support needs with any XSEDE staff? If so, please indicate with whom

Please enter your answer below: Pietro Cicotti & Bob Sinkovits

Have you received XSEDE advanced support in the past?. If so, please indicate the time period, and how the support you received then relates to the support you request now.

Please enter your answer below: We received ECSS support as a part of the COSMIC2 science gateway through Andres Leschziner's XSEDE account.

New **Approved** Opportunity Information

Resources

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

Approved	50,000.00 GB	Comments: (none)
Recommended	50,000.00 GB	Comments: (none)
Requested	50,000.00 GB	Comments: As a science gateway, the COSMIC2 gateway will support many users, each of which will have 10 - 300 GB dataset sizes. While storage will not be permanent, we are requesting to ensure that the users jobs will not be limited by storage.

2. SDSC Dell Cluster with Intel Haswell Processors (Comet)

Approved	50,000.00 SUs	Comments: (none)
Recommended	50,000.00 SUs	Comments: (none)
Requested	50,000.00 SUs	Comments: We are also requesting CPU-based SUs on Comet (in addition to Comet GPU) to ensure that we can access an older version of Relion (Relion-1.4) that was previously written as CPU-only software. Now, it is capable of GPU-acceleration, which is why we will plan on using most of our SUs on Comet GPU.
Do you need Virtual Clusters?. More about Virtual Clusters		
No		

3. Stanford Cray CS-Storm K80 (XStream)

Approved	1,000.00 SUs	Comments: (none)
Recommended	1,000.00 SUs	Comments: (none)
Requested	1,000.00 SUs	Comments: Long term, we imagine that users on our science gateway will be able to submit Relion-2.0 jobs to any GPU cluster supported by XSEDE. Therefore, we are requesting 1,000 SUs on XStream to give us the SUs required to test integration into our COSMIC2 gateway.

4. SDSC Comet GPU Nodes (Comet GPU)

Approved	26,400.00 SUs	Comments: (none)
Recommended	26,400.00 SUs	Comments: (none)
Requested	26,400.00 SUs	Comments: We are requesting 26,400 SUs for Comet to be used on the GPU nodes using Relion-2.0, a GPU-accelerated software for cryo-EM structure determination. Previously, Comet used a CPU-

based version of Relion (Relion-1.4) that was successful in calculating structures. We will be using this allocation for the computing tasks from the science gateway COSMIC2. COSMIC2 is an ECSS and SGCI-supported science gateway that will open to outside users during this submission period. We have spent the past year using the remaining SUs of an XSEDE allocation to support our job testing and now we are ready to allow many users access to this resource.

Documents

1. *Main Document*: **Main document** XRAC_June2017.pdf (958.8 kB)
2. *Code Perf & Scaling*: **Code performance and scaling** Code Performance and Scaling.pdf (19.2 kB)
3. *PI CV (2 pages limit)*: **CV** BiographicalSketch_NSF_April2017.pdf (32.9 kB)

Grants

Publications Resulting from XSEDE Work

No publications entered!

Summary	Systems Monitor	Get Started	Announcements	Overview	Forums	Overview	ECSS Overview
Allocations/Usage	Remote	Community	Resource Info	Course Calendar		Help Desk	ECSS Projects
Accounts	Visualization	Codes	Startup	Online Training		Security Incident	ECSS Symposium
Jobs	Storage	Manage Data	Education				ECSS Workflows
Profile	Software	User Guides	Research				ECSS
Publications	Queue Prediction	News	Submit/Review				Justification
Tickets	SU Converter	Usage Policy	Request				
Change	Science	Knowledge Base	Manage				
Password	Gateways		Allocations				
Add User	Scheduled		Policies				
Community	Downtimes		About XRAS				
Accounts							
SSH Terminal							

ABOUT

Welcome

Team

XSEDE Home

The Extreme Science and Engineering Discovery Environment (XSEDE) is supported by the National Science Foundation.
 For general questions, contact info@xsede.org | For user assistance, please submit a consulting ticket | ©2011 XSEDE. All Rights Reserved.