

OOKAMI

The Ookami Apollo80 system: Progress, Challenges and Next Steps

Eva Siegmann, Robert Harrison
Stony Brook University



Stony Brook
University



iACS INSTITUTE FOR ADVANCED
COMPUTATIONAL SCIENCE



University at Buffalo
The State University of New York

Fugaku

#1 Fastest computer in the world until 6/2022





First machine to be fastest in all 5 major benchmarks

- Green-500 benchmark
- Top-500 benchmark
- HPCG benchmark
- HPL-AI benchmark
- Graph-500 benchmark



- 432 racks
- 158,976 nodes
- 7,630,848 cores
- 440 PF/s dp (880 sp; 1,760 hp)
- 32 Gbyte memory per node
- 1 Tbyte/s memory bandwidth/node
- Tofu-2 interconnect

Ookami



Processor	A64FX
Memory	32GB@ 1TB/s
#Nodes	176
Comms	IB HDR-100

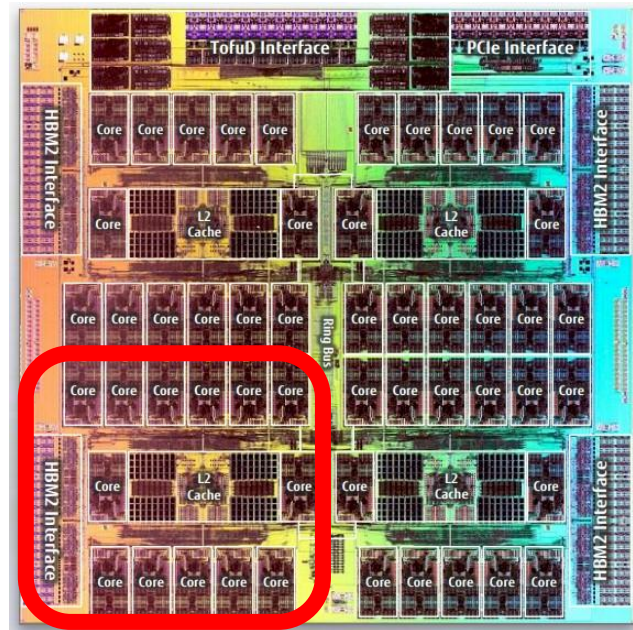


1.5 mil node hours per year

A64FX NUMA Node Architecture



- Arm V8-64bit
- Supports high calculation performance and low power consumption
- 32 (4x8) GB HBM @ 1TB/s
- Supports Scalable Vector Extensions (SVE) with 512-bit vector length
- 4 Core Memory Groups (CMGs)



Ookami Environment




- Rocky Linux 8
- Lustre file system providing ~800TB
- Slurm workload manager
- Full bisection bandwidth HDR100
- Module environment
- Various compilers, profilers & debuggers

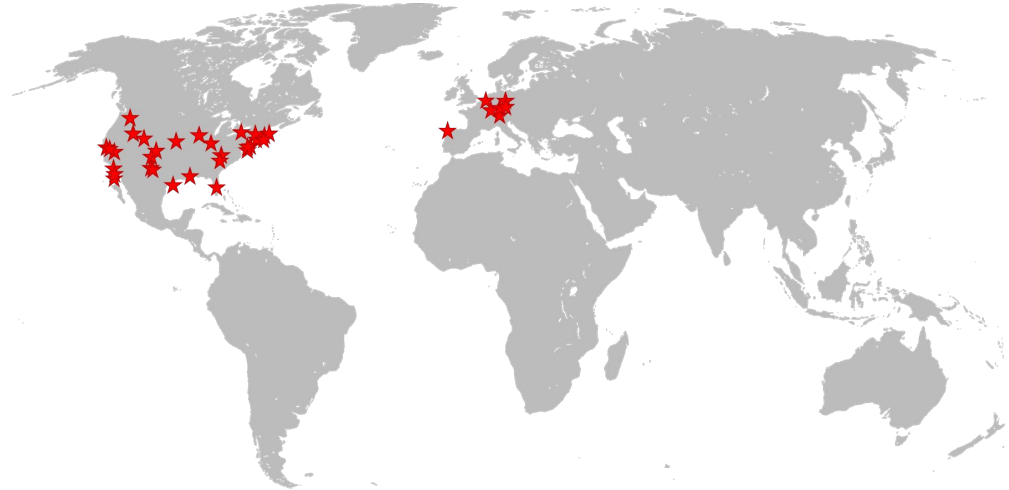
User Community (April 2023)



2019 2020 2021 2022 2023 2024 2025



- 105 projects
- 96.6% academia
- 3.4% industry (using so far 0.095% of all available node hours, no industry production projects)
- 90% USA, 10% Europe
- 300 users



User Support



- Website
- Ticketing system
- Slack channel
- Office hours
 - Virtual meetings twice a week
(Tue 10am-noon, Th 2pm-4pm)



The screenshot shows the 'Open a New Ticket' page of the iACS Support Center. At the top, there is a navigation bar with 'Support Center Home', 'Open a New Ticket', and 'Check Ticket Status'. The main heading is 'Open a New Ticket' with a sub-instruction: 'Please fill in the form below to open a new ticket.' Below this is a section titled 'Open a New Ticket' with a paragraph of instructions: 'If you would like to open a ticket to request a project number, request an account, report a problem, give us feedback, or if you have a general inquiry, please fill out the form below. Choose the topic of your ticket from the Help drop down list. Once the information is received, you will be issued a ticket number. If you would like to access a ticket already assigned a number, click on the Check Ticket Status at the top of the form.' The form fields include: 'Name *' (text input), 'Email Address *' (text input), 'Phone Number' (text input) and 'Ext.' (text input), 'Net ID. If no net ID, create a user name *' (text input), 'Help Topic' (dropdown menu), and 'CAPTCHA Text:' (text input) with a pink box containing '339F9' and a small image. At the bottom right, there are three buttons: 'Create Ticket', 'Reset', and 'Cancel'.

Webinars & Courses



- Webinars about profilers, compilers, SVE hackathons, etc.
 - All recordings and slides can be found on our [webpage](#)
- Webinars targeting potential new users (e.g. Campus Champions, conferences)
 - Close cooperation with organisations like the Arm HPC User Group, Cray User Group, ACCESS, etc.
- A64FX symposium
 - Organized and initiated by Ookami team
 - Speakers from SBU, UDEL, LANL, University Bristol UK, NHR@FAU Germany, Total Energies, Arm, HPE, Fujitsu
- Annual Ookami user group meeting

Around 70 presentations and events

Events - Introductory HPC Courses for High School Students




- Several courses (2 days each)
- Organized in cooperation with:
 - CSTEP/STEP (Collegiate Science and Technology Entry Program)
 - Institute for Electrical and Computer Engineering
- Evaluation survey before and after the course



Evaluation - HPC for High School Students




How confident do you feel about HPC?




What did you enjoy the most?


- Visiting the server room
- Seeing all the applications using HPC
- Working on the cluster
- Making new friends
- Doing the puzzle
- I liked the astrophysics talk a lot, because it was a real life example of how what we're learning can be applied, and the speaker was very passionate about her subject.



How confident are you using the command line?



Are you excited about the course / Did you enjoy it?



- *I loved how the instructors helped everyone out before actually moving to the subject.*
- *I wish there was more group activities like the puzzle, but the class was great.*
- *This introduction was very informative and easy to follow. I appreciate having taken it!!*

Funding Grad Students



- Increasing the number of science codes effectively using A64FX
- Broadening the number of disciplines
- Accelerating adoption of the system for science including partnering with application teams to develop successful ACCESS applications
- Transferring knowledge into the science teams and communities
- Translating success on Ookami into performance and productivity advances on other NSF cyberinfrastructure employing mainstream x86 technologies

Science Codes Ported by Graduate Students

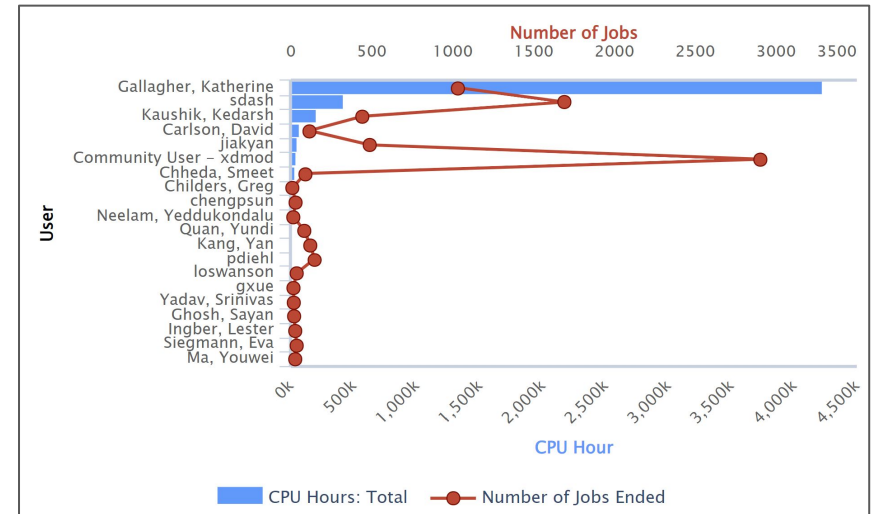


Code	Field
FLASH	Astrophysics
MOM6	Ocean Model
siesta	Electronic structure and dynamics
amber	Biomolecular simulation
pytorch	Machine learning
msieve	Library for factoring large integers
BUFIA	Linguistics application

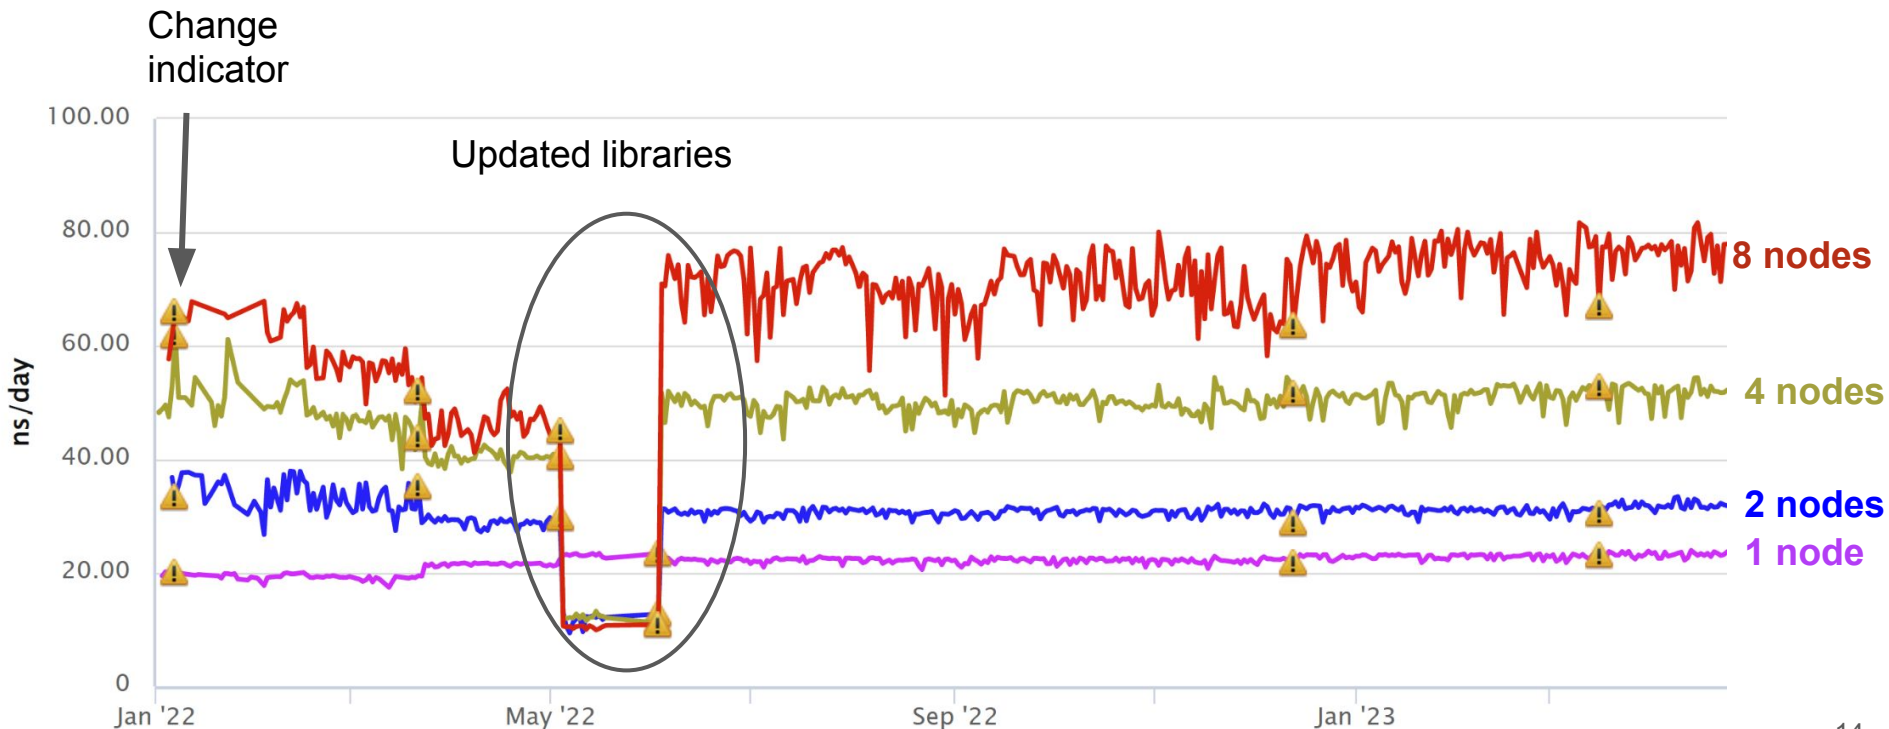
XDMoD: A Comprehensive Tool for HPC System Management



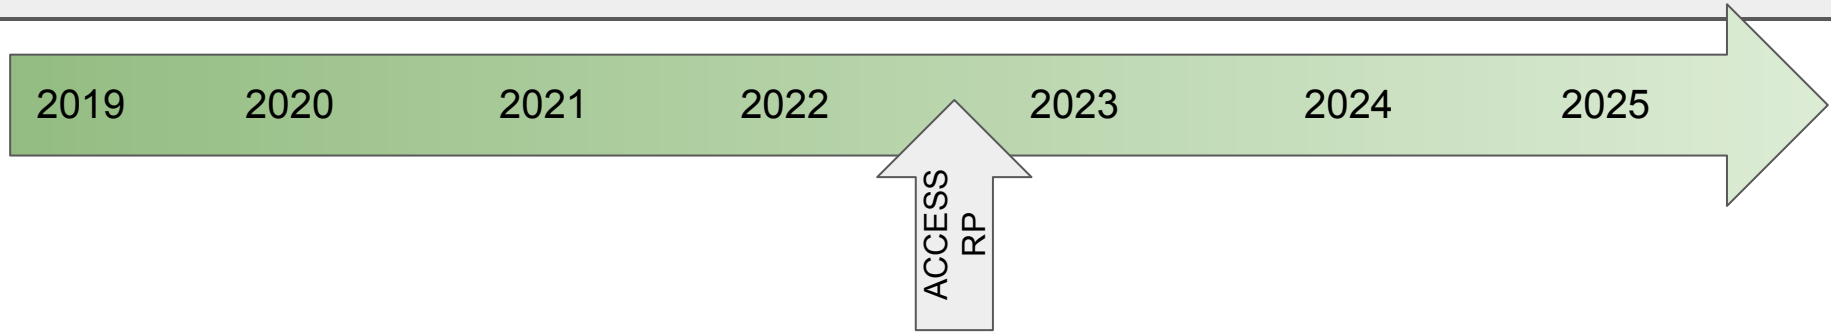
- Optimize Resource Utilization and Performance
 - Provide detailed information on utilization
 - Measure and improve job and system level performance
- Detailed performance data collected from nodes, e.g.:
 - SVE instruction count
 - Power consumption
 - Memory usage



Performance over Time: GROMACS



ACCESS Integration

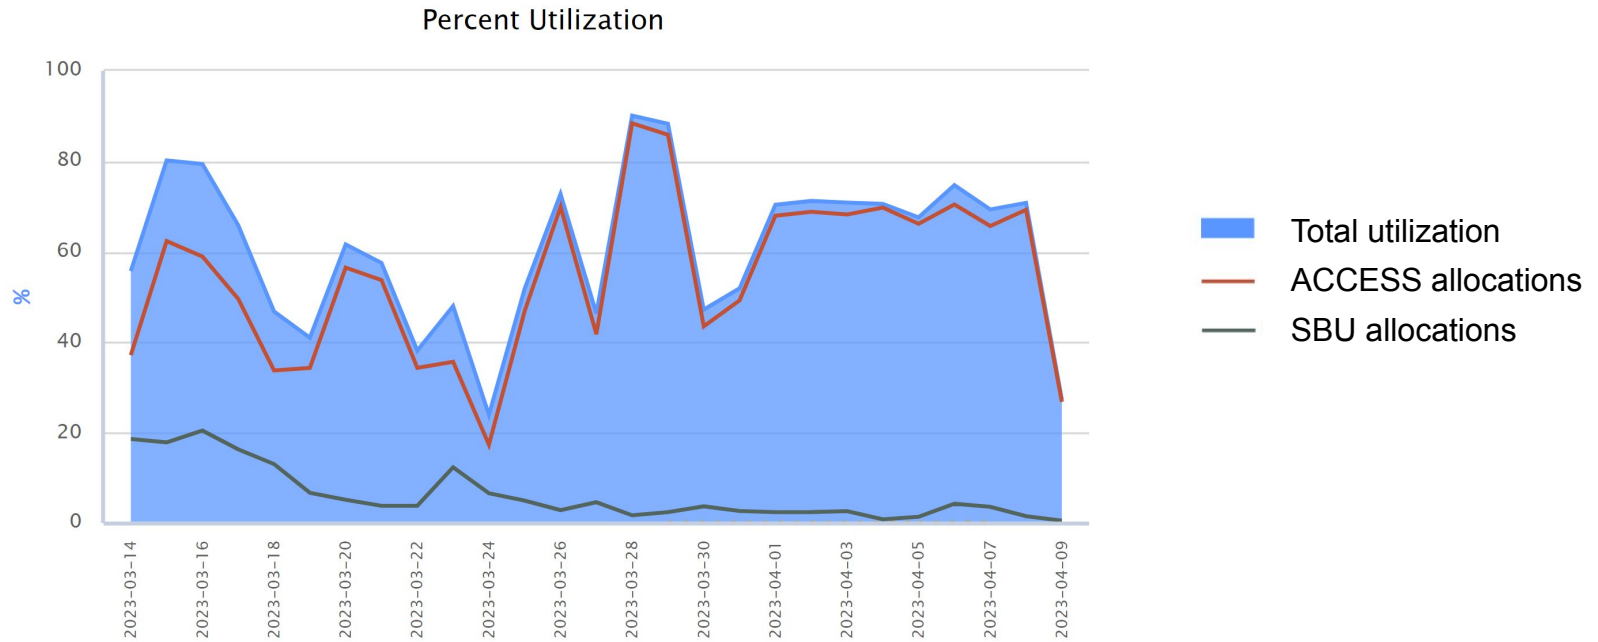


- All requirements successfully fulfilled
- Actively collaborating with the ACCESS community
- Ookami is actively fulfilling (exceeding!) its expectations as new ACCESS Resource Provider

ACCESS Projects



- 13 ACCESS projects





Use Case:

Modeling the Physical Oceanography
Adjacent to Pygoscelis Penguin Colonies
along the West Antarctic Peninsula

Catherine Gallagher, SBU

Penguins!

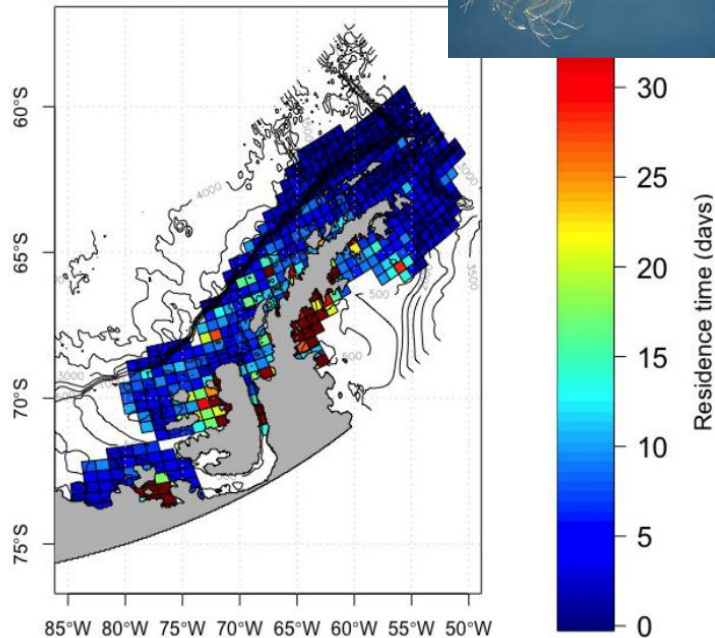


- Global warming is huge concern
- Investigating *Pygoscelis* penguins
- Relationship between penguin diets, local prey retention, and colony presence

Simulation



Particles released throughout the peninsula
simulating krill behavior



ROMS

Regional Ocean Modeling System


- Free-surface, terrain-following, primitive equations ocean model
- Open-source
- Fortran
- Ported and optimized for A64FX

Performance Comparison



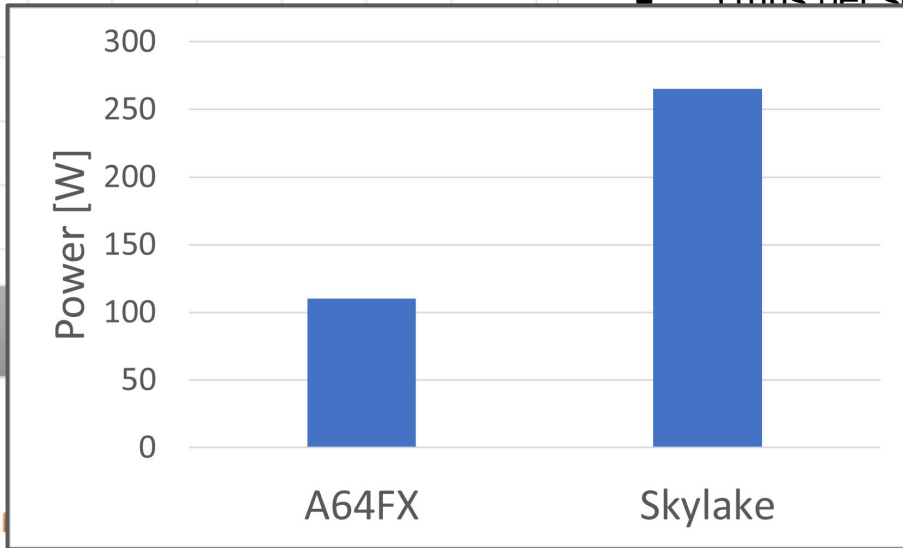
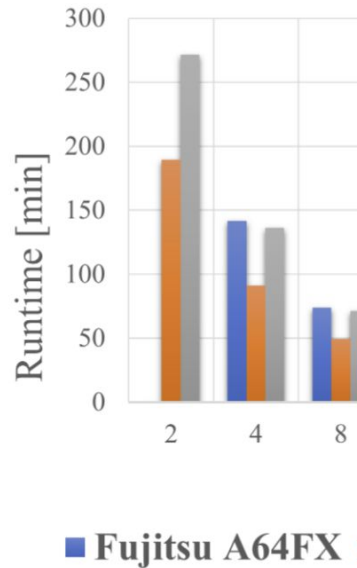
- Intel Skylake nodes (Xeon Gold 6148 CPU), 2.40GHz, 40 cores
- Intel Haswell nodes (Xeon CPU E5-2690 v3), 2.60Ghz, 24 cores
- High-speed InfiniBand FDR network
- Same versions of
 - GCC
 - OpenMPI
 - NetCDF

Performance Comparison



- 5 runs per setup - deviation of max 2.5%
- Performance parity at 50 nodes
- Speculation - A64FX might have better performance than Intel nodes at higher node counts

Performance Comparison



- 5 runs per setup - deviation of max 2.5%
- parity at 50 nodes
- A64FX might have better than Intel nodes at higher

Lessons learned



- Users need extensive support to efficiently use the system
- Compiler plays a major role in achieving good performance
- Most science codes need optimization to fully utilize A64FX and give the best performance
- Vectorization and threading is crucial

Future goals



2019 2020 2021 2022 2023 2024 2025

End

- Increase utilization and especially increase ACCESS allocations
- Continue to partner with ACCESS and serve the national community
- Advance testbed projects to production projects
- Publications especially joint with users
- Expand and enhance educational initiatives & community facing activities
- Support the success of grad students' science projects; advance to production projects

The Ookami Team



University at Buffalo

Matt Jones (Co-PI)
Nikolay Simakov
Joseph White



System architect: Firat Coşkun
Admin: Daniel Wood
Admin: John Dey

Chief Research Information
Officer: David Cyrille



SBU Scientific SW support

Tony Curtis
Eva Siegmann
Dave Carlson



SBU graduate students

Catherine Feldman (Astro.)
Smeet Chheda (CS)
Kedarsh Kaushik (Physics)
Youwei Ma (Marine & Atm. Sciences)
Yuzhang Wang (Phys. & Quant. Biology)
George Liang (HPC support)
Logan Swanson (Linguistics)
Chengpeng Sun (Appl. Mathematics & Statistics)



PI: Robert Harrison
Co-PI: Barbara Chapman
Co-PI: Alan Calder



www.stonybrook.edu/ookami/

ookami_computer@stonybrook.edu

Thanks to the whole project team and the NSF!