

Exploring GPUs for Environmental Omics Analysis

Kjiersten Fagnan, PhD
CIO, Data Science and Informatics Leader
kmfagnan@lbl.gov



Talk Outline



- What is the DOE Joint Genome Institute?
- Overview of 'omics data generation
- Data and analysis challenges in genomics
 - Scaling
 - Exascale computation
 - Hackathons to bring the community together

Outline





The JGI is a DOE Office of Science User Facility



Each Office of Science user facility exists through investment by a program "owner" – the Office of Science program that provides funds through congressional appropriations for construction and operations. The decision to invest in a user facility emerges through long term strategic planning with the scientific community to identify the research tools that will deliver the greatest scientific impact to advance the DOE mission.

Each user facility represents a substantial commitment on the part of its sponsoring program, which provides oversight and works closely with the facility management to maximize scientific impact and productivity.



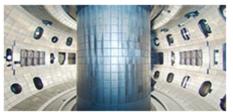
ASCR User Facilities



BES User Facilities



BER User Facilities



FES User Facilities



HEP User Facilities



NP User Facilities

A User facility is a federally sponsored research facility available for external use to advance scientific or technical knowledge

DOE Joint Genome Institute



Who we are...



JGI serves the diverse scientific community as a national user facility, enabling the application of large-scale genomics and analysis of plants, microbes, and communities of microbes relevant to the DOE mission goals of harnessing science and technology for addressing energy and environmental challenges.

A U.S. Department of Energy Office of Science User Facility

- Walnut Creek, CA facility opened in 1999
- ~280 staff;
 - ~\$69M annual funding















The DOE JGI Continues to Evolve



2004 2012 2017.... 1997 Structure The Human A Production A Next-Generation An Integrative Genome Science Genome Sequencing Genome Science **User Facility User Facility** Project User Facility nature Human & Plants & Host Interactions. **INTEGRATIVE** Animals Microbes Pathway Modeling, **OMICS** DNA Synthesis, Science Metabolomics **Facilities Integrating** PNAS Science nature Collaborations for User Science **KBase** ROOT DWELLERS

lission

sequence, assemble and annotate **GENOMES**

understand gene and genome FUNCTION

harness sequence and functional knowledge for **SOLUTIONS**

DOE Systems Biology Knowledgebase

JGI's Capabilities are Evolving to Meet **Diverse User Needs**



Infrastructure

Bioenergy Carbon Cycling Mission Areas Plants Fungi Metagenomes **Microbes Programs DNA Sequencing Advanced Genomic**

- Illumina platforms: NovaSeq, HiSeq, MiSeq, NextSeq
- PacBio platforms: RSII, Sequel
- Oxford Nanopore platforms: MinION, PromethION

Technologies

- Single-cell Seq
- Metabolomics
- Methylation
- Epigenomics
- Transcriptomics



Computational Analysis

- Assembly/ Annotation/Analysis
- High-Performance Computing
- Big Data Integration & **Analysis Systems**



DNA Synthesis Science

DNA Synthesis

Biogeochemistry

- Design & Pathway Assembly
- Host Engineering

Serving a Global User Community

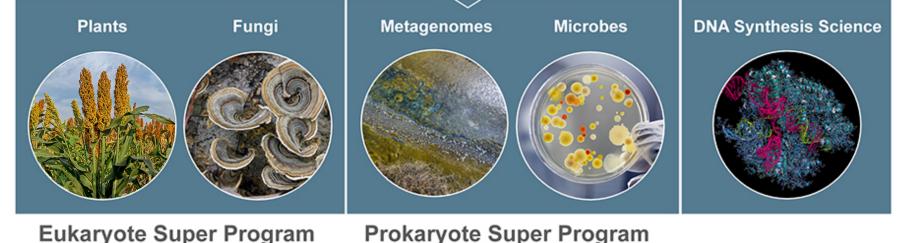




User Programs and Science Programs

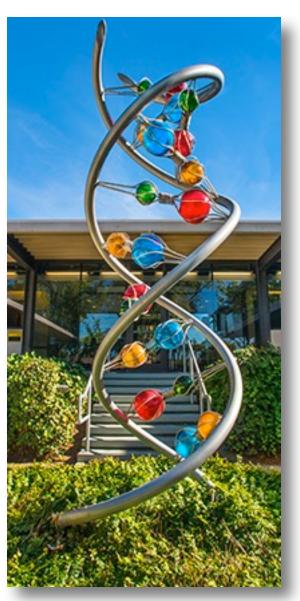




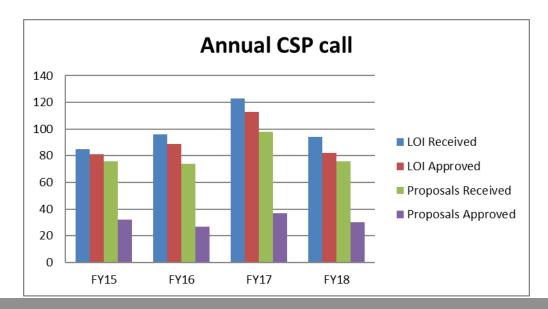


Large-scale Community Science Program





- Flagship user program established in 2004
- Current focus is on large-scale, multiinvestigator projects in specific theme areas
 - Plant functional genomics and microbiomes
 - Function-driven microbial genomics
 - Microbial communities in nutrient cycling
 - Extreme environments
- All capabilities offered; 2 Tbp sequence cap



New Sequencing Approaches for Identifying Novel Organisms



New Sequencing Efforts

GEBA Phases II and III

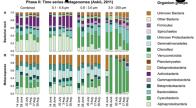


CSP in progress
10,000 reference genomes?

Algal Diversity



Temporal/Spatial Metagenomes



New Sequencing Technologies



NovaSeq



PromethION

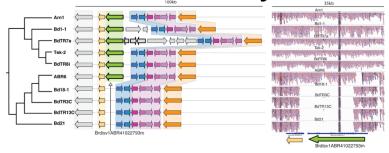
- Increased output (300 Tb)
- Real-time/rapid output
- Metagenomes/transcriptomes
- Epigenomes
- Single cells
- Functional genomics

Future Directions: Leveraging New Analysis Techniques

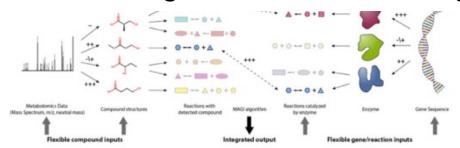


New Analysis Tools and Algorithms

Pan Genome Analysis



Metabolites, Annotation and Gene Integration consensus scoring



Microbiome Data Science



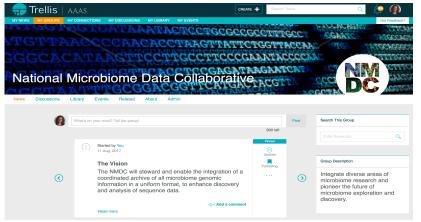




Future Directions: Multi-omics Microbiome Data Analyses









National Microbiome Data Collaborative

Enable an integrated microbiome data ecosystem across diverse areas of research in medicine, agriculture, bioenergy, and the environment.

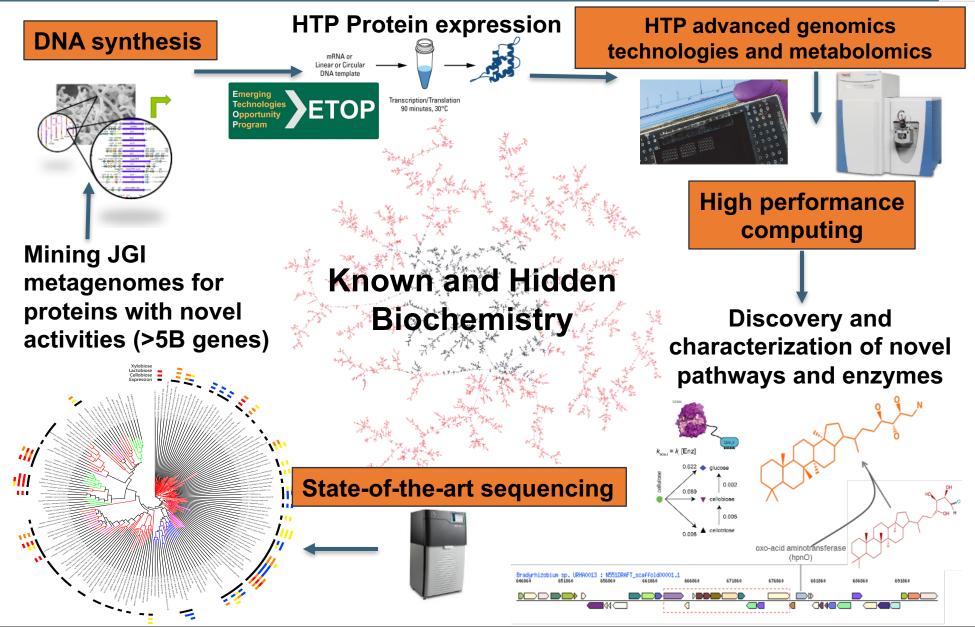






Future Directions: Linking Sequence Data to Function

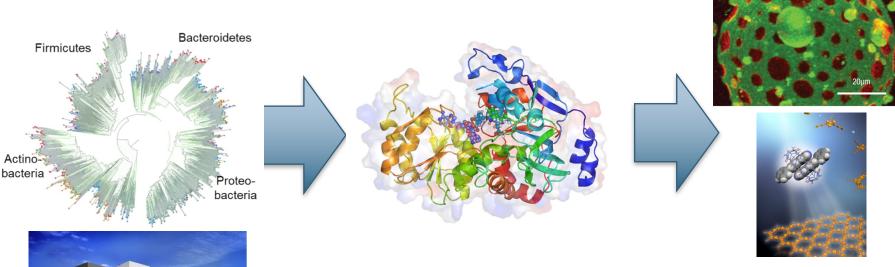




Future Directions – Integrating Diverse Data Types



From Genes to Structure to Nanomaterials





















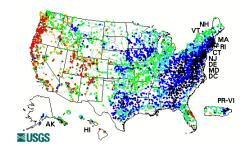


Future Directions - Combining experimental data with large scale simulation to understand watershed dynamics











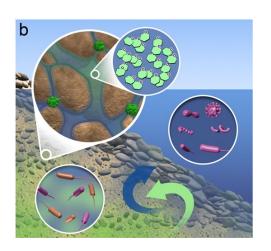


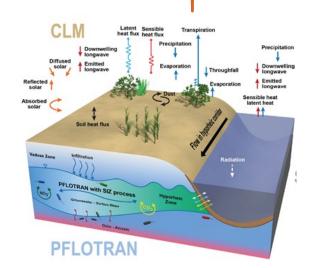


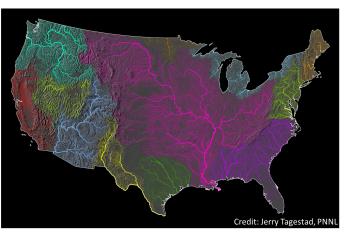




$$r_i = \mu^{\text{max}} \exp\left(\frac{-Y_{Cs,i}}{V_h[Cs_i]}\right) \exp\left(\frac{-Y_{O_2,i}}{V_h[O_2]}\right), \quad i = 1, 2, ..., N$$







Outline

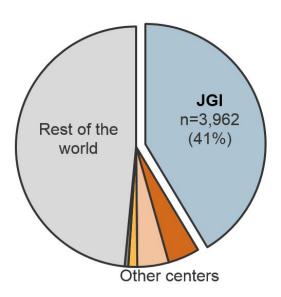




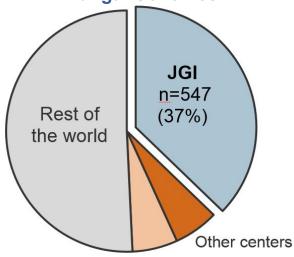
JGI is a World Leader in Genome Production



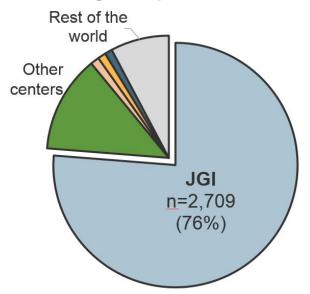
Bacterial and Archaeal Isolate Species



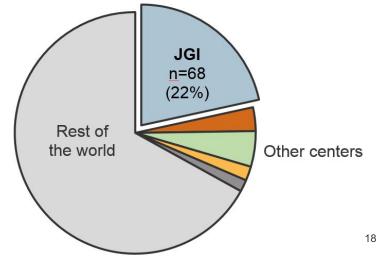
Fungal Genomes



Bacterial and Archaeal Single-Amplified Genomes

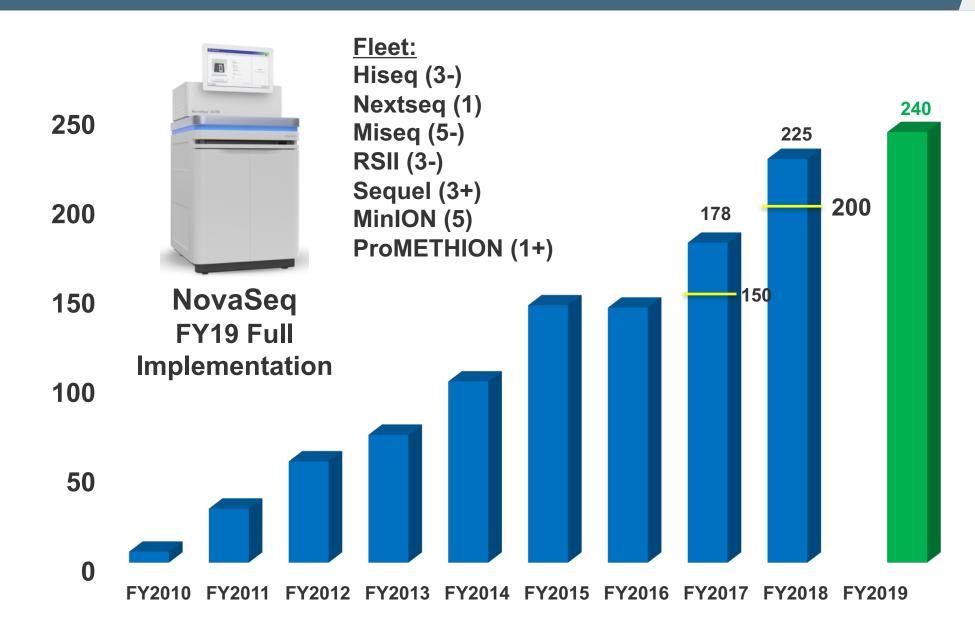


Plant Genomes



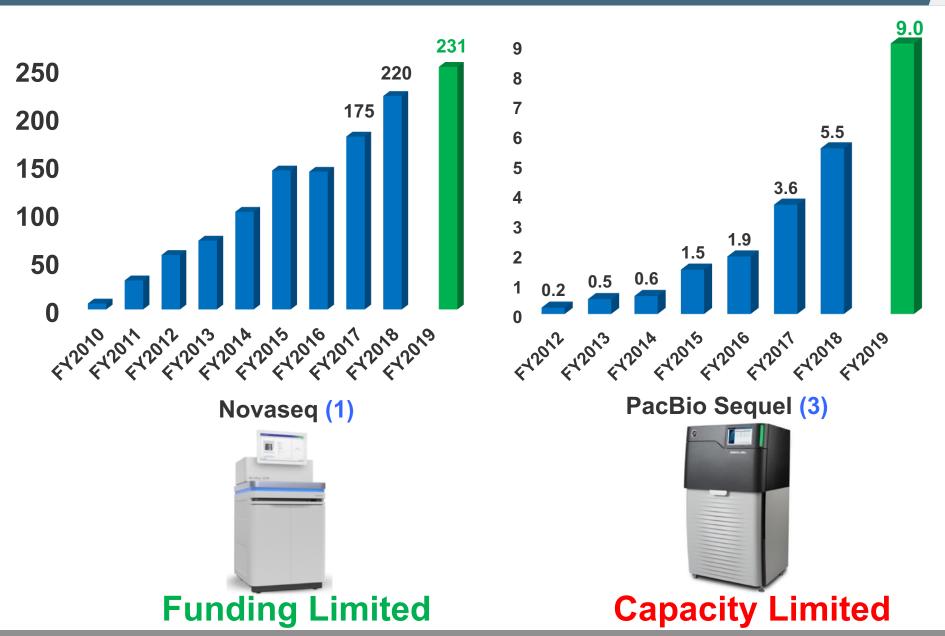
JGI Yearly Base Output (Tb)





Short versus Long Read Platforms

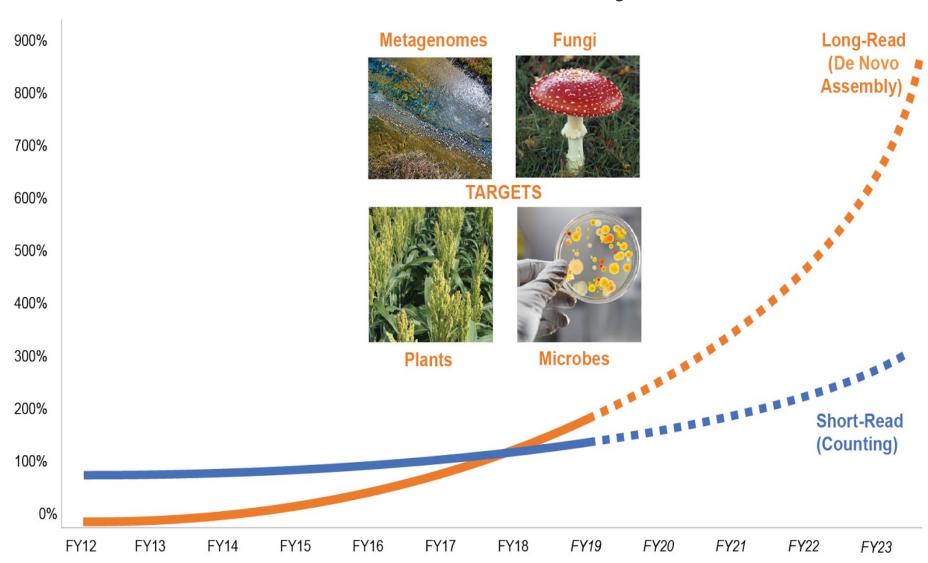




Short versus Long Read Platforms



JGI Relative Growth Projection

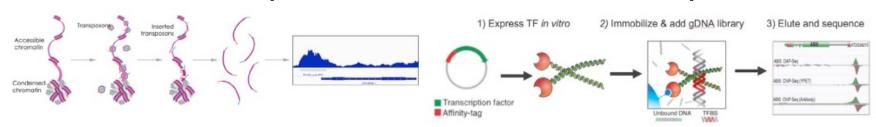


JGI Continues to Develop in State-of-the-Art Technologies

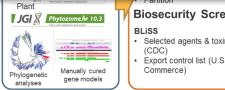


Functional Genomics ATAC-Seq

DAP-Seq



New data types require new analysis techniques















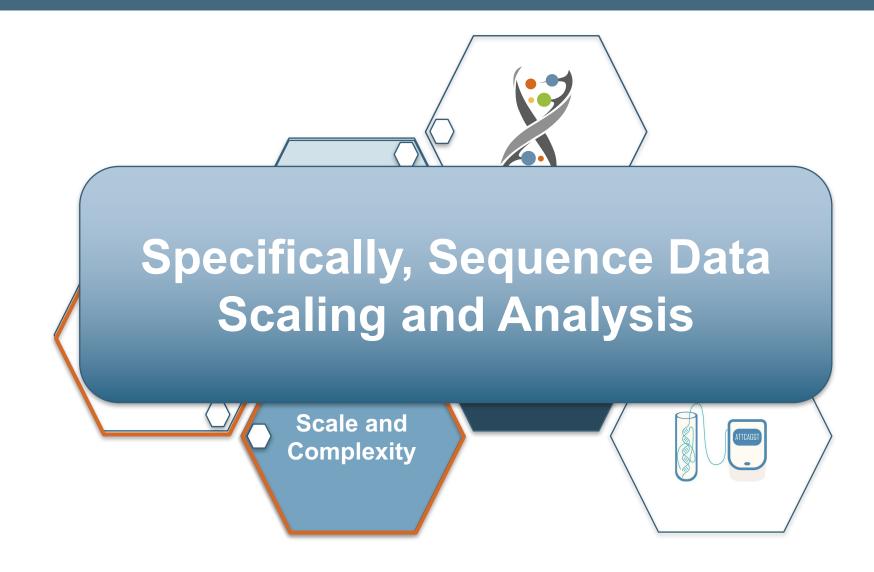
Yasuo Yoshikuni



Trent Northen

Outline

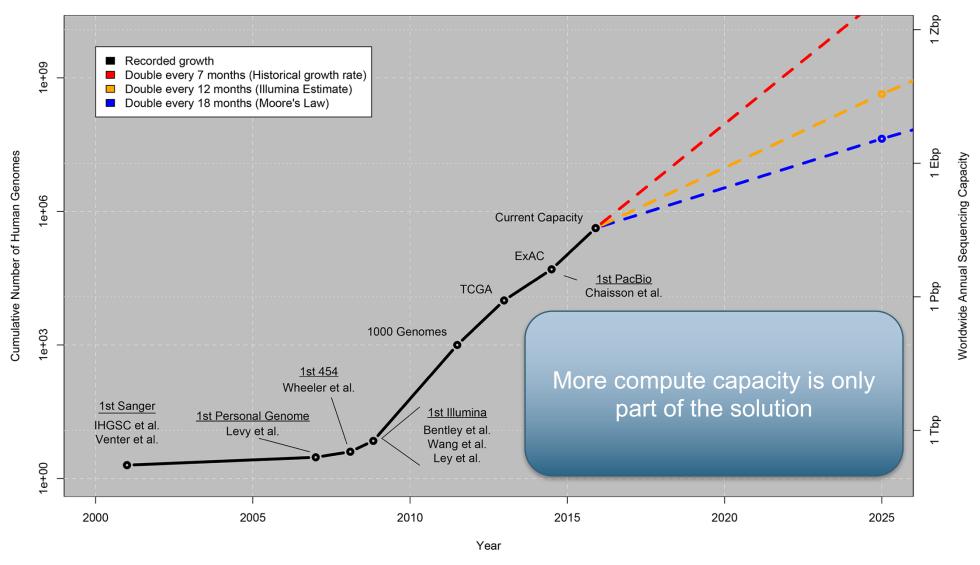




Sequence generation is outpacing Moore's law



Growth of DNA Sequencing



Bioinformatics computing needs have evolved



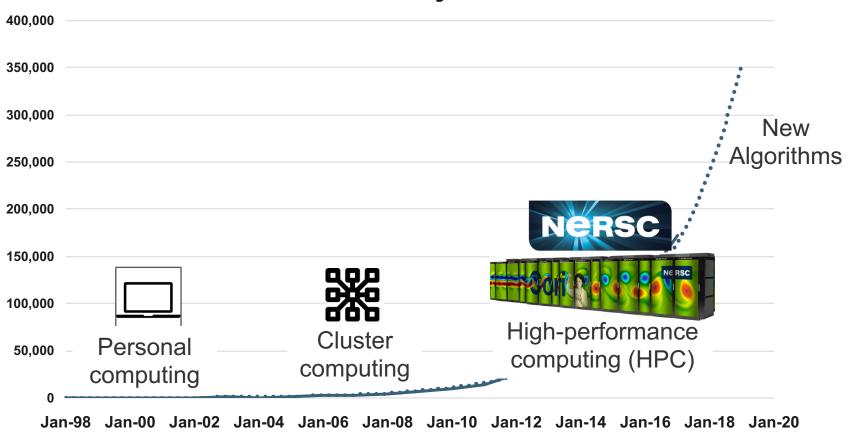
NVIDIA®

'intel

inside"

XEON PHI

Total Genome Projects in GOLD



25

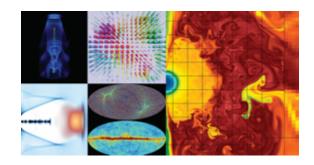
Source: IMG-GOLD statistics portal

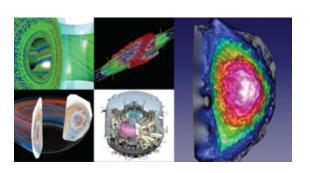
NERSC is the mission HPC computing center for the DOE Office of Science



- NERSC is a DOE User Facility that deploys advanced High-performance Computing and data systems for the broad Office of Science community
- Approximately 6000 users and 750 projects
- Over 2000 publications resulting from NERSC resources per year
- Supports many other DOE User Facilities including the DOE JGI







JGI-NERSC Partnership

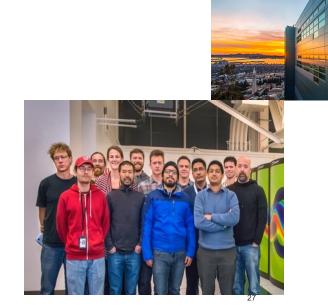


Started in 2010 when JGI needs outpaced in-house capacity



 NERSC procures, manages and troubleshoots JGI's compute infrastructure

- JGI and NERSC staff partner on computing initiatives
- NERSC's Cori supercomputer is comprised of Haswell and Xeon Phi hardware



NERSC Deploying GPU-powered System



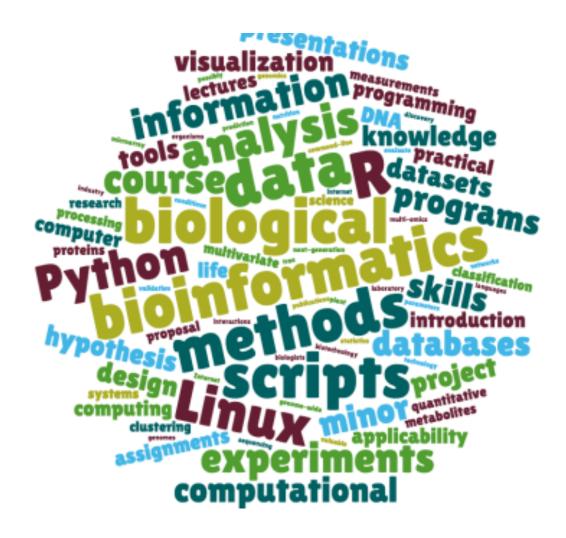
Perlmutter: A System Optimized for Science

- Cray Shasta System providing 3-4x capability of Cori system
- First NERSC system designed to meet needs of both large scale simulation and data analysis from experimental facilities
 - Includes both NVIDIA GPU-accelerated and AMD CPU-only nodes
 - All-Flash filesystem for I/O acceleration
- Robust readiness program for simulation, data and learning applications and complex workflows
- Delivery in late 2020



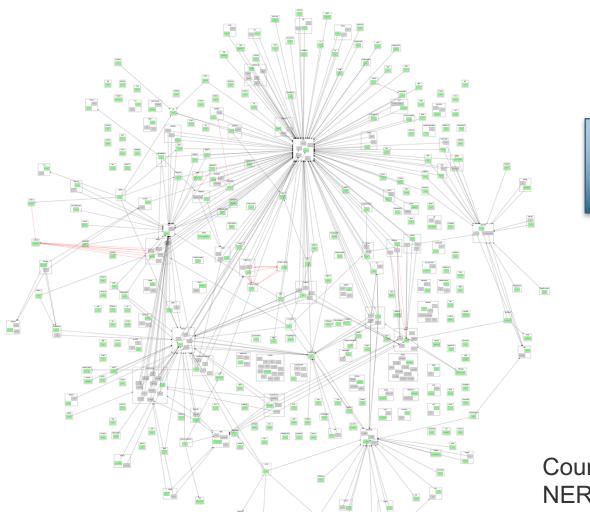
Optimizing the JGI Workload for Perlmutter





Software Supporting JGI Analyses is Diverse





More than 2000 software packages and modules in use

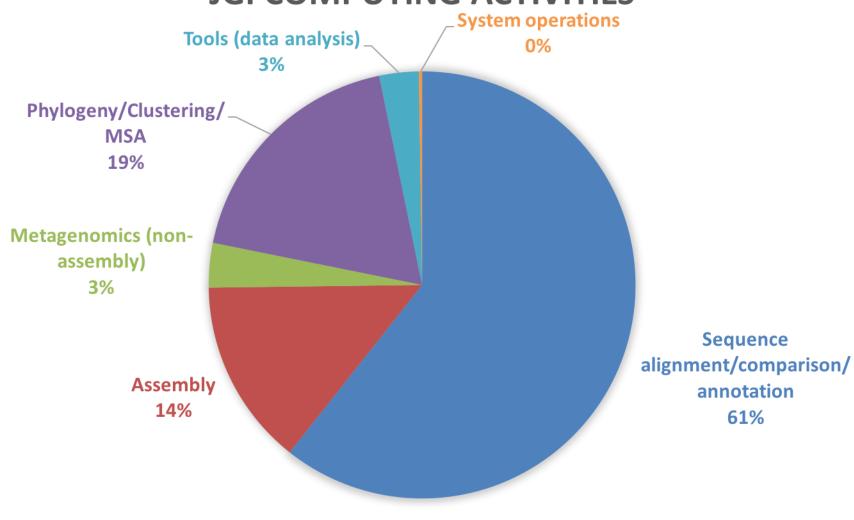


Courtesy of Doug Jacobsen, NERSC

Most Compute Time Used for Sequence Comparison





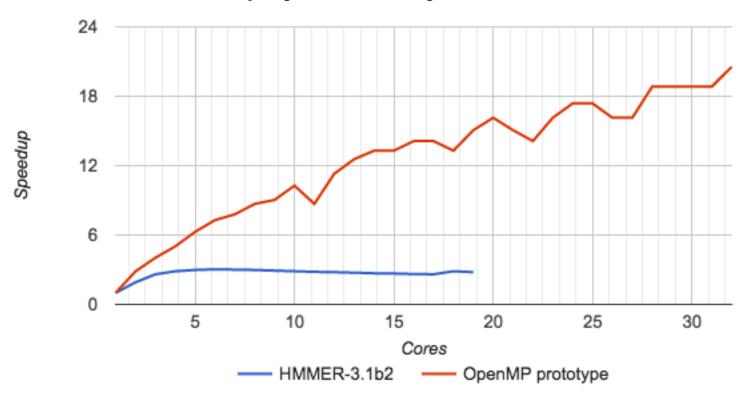


Single Node Optimization of Sequence Comparison – Xeon Phi



Adding OpenMP to hmmsearch led to improved scalability

Deployed on JGI systems in 2016



Courtesy of Bill Arndt, Postdoc - NERSC Funded by LBNL Intel Parallel Computing Center

Multiple Node Optimization of Sequence Comparison with Spark



BLAST is ~41% of the JGI computing workload

 BLAST is pleasingly parallel – can be broken into independent chunks - what's the best way to distribute the work?



Widely-used tool for distributing compute tasks

SparkBLAST implementation by *Chris Beecroft*, Data Management Group, Genomic Technologies

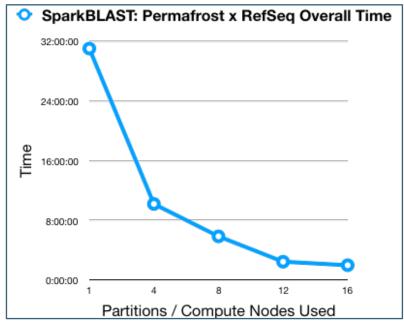
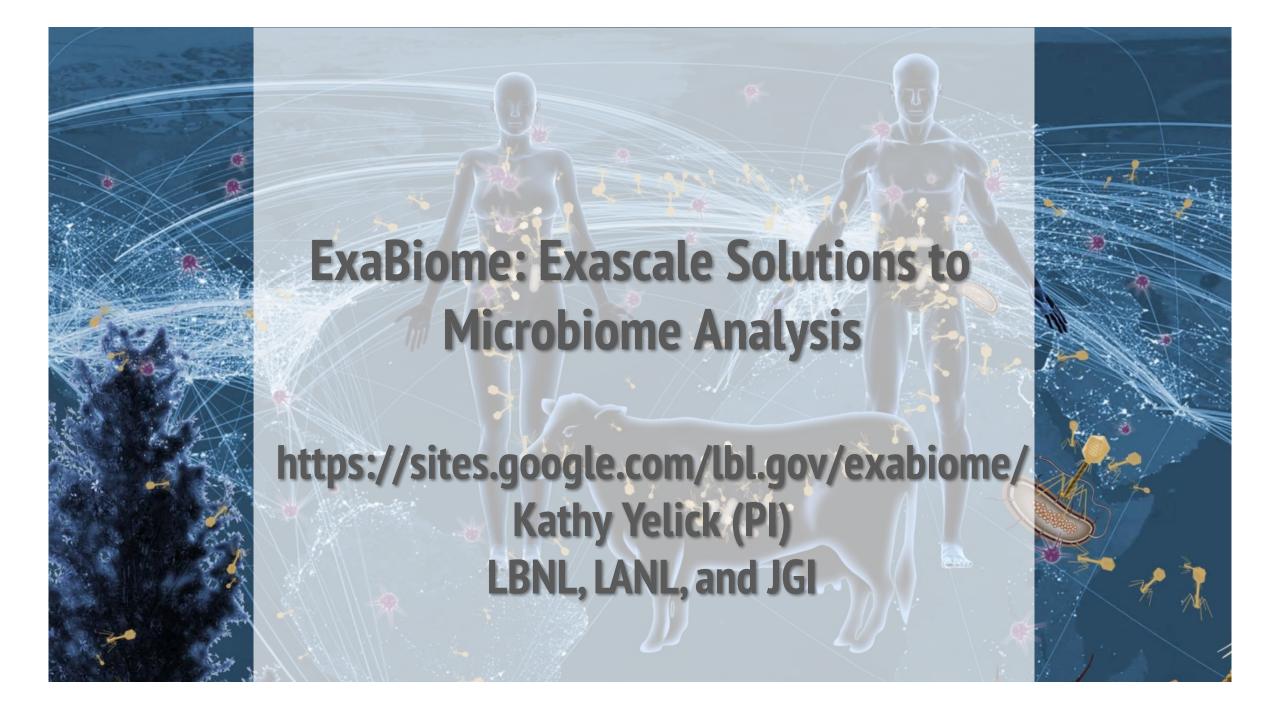


Figure: SparkBLAST is ~15X faster on 16 nodes than on a single core enabling significantly higher throughput



ExaBiome: Exascale Solutions for Microbiome Analysis



- Microbes: single cell organisms, such as bacteria and viruses
- Microbiomes: communities of 1000s of microbial species, less than 1% individually culturable in a lab (and thus sequenced)
- Metagenomics: genome sequencing on these communities (growing exponentially)



Environment



Health



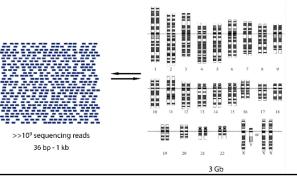
Agriculture

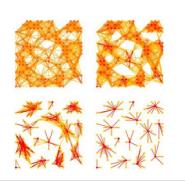


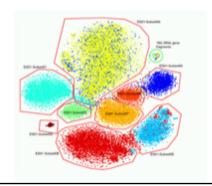
Bio-Manufacturing

ExaBiome: Exascale Solutions for Microbiome Analysis









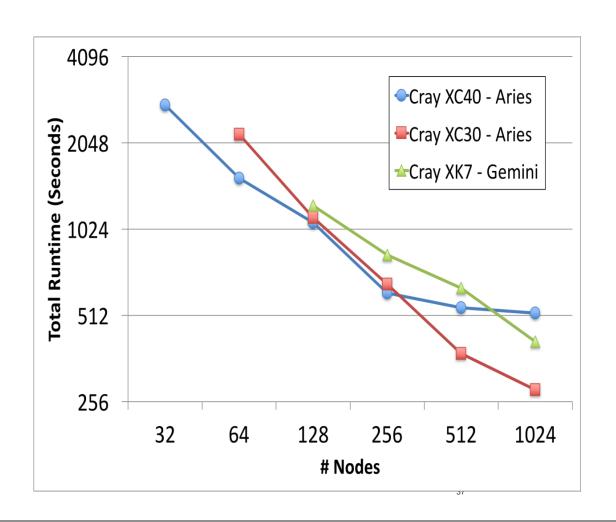
Metagenome Assembly	Protein Clustering	Comparative Analysis
Graph algorithms, Hash Tables, alignment (Smith-Waterman)	Machine learning (clustering), sparse linear algebra / graphs	Alignment, Machine learning (dimensionality reduction), linear algebra
Fine-grained comm., all-to-all, remote atomics and fast I/O	Fast barriers, subset reductions	All-to-all

Multi-Node Strong Scaling – Genome Assembly



HipMer scales to a thousand nodes (10Ks of cores) on a fixed modest sized problem (human genome)

- Programming model lets HPC system "look like" shared memory
- HPC network is critical: on a 10G ethernet cluster, this problem is 10x slower on 2 nodes than on 1



Exascale Bioinformatics Computing on Summit



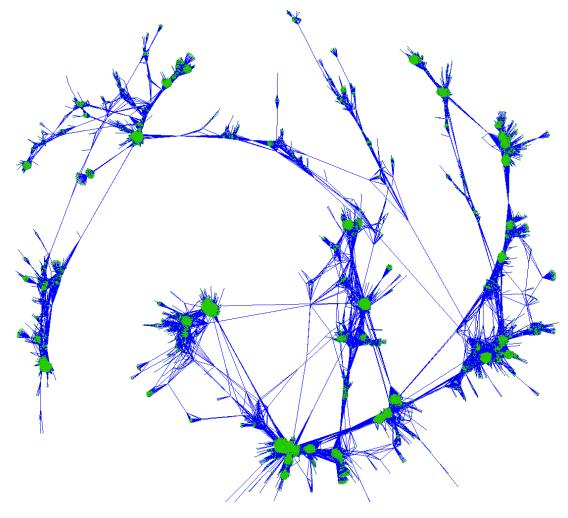




Co-evolution: Big data and high performance computing reveals underlying biological signatures



- 10 million SNPs
- >1 x 10¹⁵ allele-specific SNP correlations calculated
- Threshold: 1 in 3 million retained
- Results modeled as a co-evolution network
- Network topology reflects the underlying biology
 - Genes under the same or similar selective pressures tend to co-evolve – which is reflected in SNP correlations and therefore, network topology
- Address population structure
 - Do this thousands of times...
 - The need for speed!



One connected component of the SNP Correlation Network

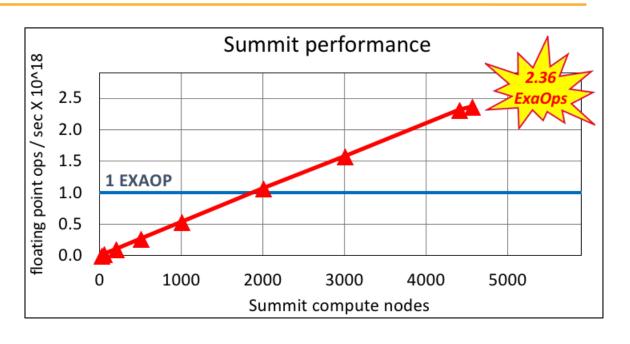


SNP Correlation/Co-evolution – Leading the Way to Exascale



Combining big data with big iron enhances statistical significance

- Thus far we have achieved <u>2.36 Exaflops₁₆</u> (mixed precision) at 4,560 nodes (99% of Summit) using the Tensor Cores
- Equivalent to 86.4 TF per GPU for the whole computation (including communications and transfers) at 4,560 nodes
- Excellent scaling made possible by Summit fat tree network with adaptive routing
- > 15,000X faster than the closest competing code



2018 Gordon Bell Prize winner

First ever for Genomics/Systems Biology

Courtesy of Dan Jacobson, ORNL

Sequence Analysis Community Slow to Adopt GPUs



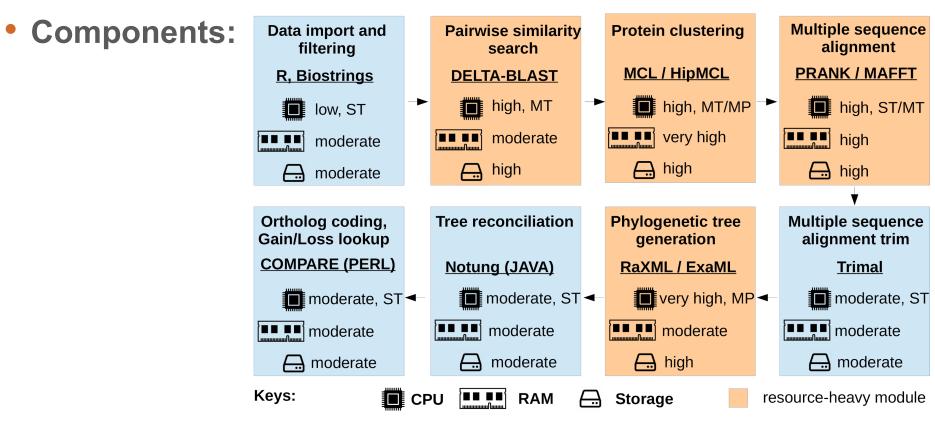
Sequence alignment				
	Tool name	Speed-up	Parallel solution	Reference
equence alignment based n BWT	BarraCUDA	-	GPU	[22]
quence alignment based BWT	CUSHAWGPU	-	GPU	[23]
quence alignment based BWT	GPU-BWT	-	GPU	[24]
quence alignment based BWT	SOAP3	-	CPU-GPU	[25]
uence alignment based hash table	SARUMAN	-	GPU	[26]
quence alignment with s based on BWT	SOAP3-dp	_	CPU-GPU	[27]
ol to map SNP ploiting SOAP3-dp	G-SNPM	_	CPU-GPU	[28]
quence alignment bloiting SOAP3-dp	G-CNV	18×	CPU-GPU	[29]
ignment of gapped short ids with Bowtie2 gorithm	nvBowtie	8×	GPU	[30]
ignment of gapped short ds with Bowtie2 orithm	MaxSSmap	-	GPU	[31]
ads assembly exploiting de Bruijn approach	GPU-Euler	5×	GPU	[32]
ads assembly exploiting de Bruijn approach	MEGAHIT	2×	GPU	[33]
quence alignment ainst database) tool	-	2×	GPU	[34]
quence alignment ainst database) tool	CUDA-BLASTP	6×	GPU	[35]
quence alignment gainst database) tool	G-BLASTN	14.8×	GPU	[36]
quence alignment with hith-Waterman method	SW#	_	GPU	[37]
quence alignment based suffix tree	MUMmerGPU 2.0	4×	GPU	[38]
quence similarity ection	GPU CAST	10×	GPU	[39]
quence similarity ection based on profiled dden Markov Models	CUDAMPF	11–37×	GPU	[40]
ultiple sequence gnment with Clustal	CUDAClustal	2×	GPU	[41]
ultiple sequence gnment with Clustal	GPU-REMuSiC	_	GPU	[42]

- There has been a lot of work to speed up similarity search codes for GPUs
 - Speedups as great as 37X
- ... but the community doesn't use these codes.
- Scientific community want to leverage solutions from the primary code developers
 - E.g. *BLAST that's not distributed by NCBI, is not acceptable to most scientists, even at JGI

JGI is Embarking on Staff Training and Collaboration



- JGI and NERSC have partnered to run a joint proposal process
 - A complex fungal comparative evolutionary analysis pipeline will be ported to Perlmutter



ST: single-threaded MT: multi-threaded MP: multi-threaded with MPI support

Hackathons to help the community



- GPU Bioinformatics Hackathon May 6-10
 - –@JGI in Walnut Creek







Come Code With Us!

Summary



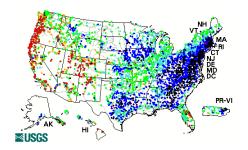
- JGI is a DOE User Facility that generates next-generation sequence data for complex environmental samples
- JGI is partnered with DOE computing facilities that provide storage and compute resources for data processing and analysis
 - These facilities are moving toward GPUs
 - More and complex biological data translates to exascale computing needs for data analysis
- Bioinformatics community uses lots of software packages, but heavily relies on core historic codes for publications
 - Community is less concerned with speed, primary concern is scientific accuracy and reprocducibility
- GPUs have been in use for several years, but have not gained traction in sequence-based analysis
 - JGI and NVIDIA hope the hackathons can help
 - Long-reads (more data) drive the need for more compute, not just IOps a better fit for accelerators

GPU applicability depends on the problem... but we have a diverse set of problems...











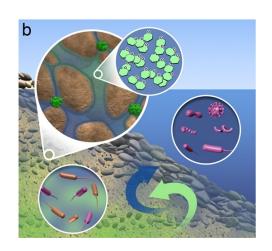


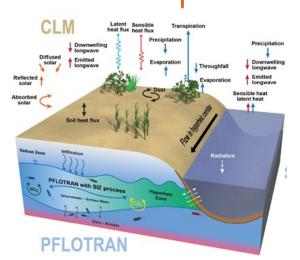


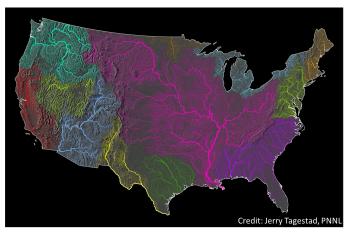




$$r_i = \mu^{\text{max}} \exp\left(\frac{-Y_{Cs,i}}{V_h[Cs_i]}\right) \exp\left(\frac{-Y_{O_2,i}}{V_h[O_2]}\right), \quad i = 1, 2, ..., N$$









Thank you!

kmfagnan@lbl.gov

