

How To Scale From Workstation Through Cloud To HPC In Cryo-EM Processing Presented by: Dr Lance Wilson

#### What will we discuss... **Presentation Outline**

- Who am I? Where am I from?
- What is cryo-electron microscopy?
- How long does analysis take?
- What impact does hardware have?
- What options do you suggest?
- What processing opportunities exist?
- Next opportunity: Light sheet microscopy!





## Who am Land Where am I from? The research context



US Dept of State Geographer © 2018 Geogle Image U.S. Geological Survey Data Sio, NoAA, U.S. Navy, NGA, Gasco

#### Google Earth

Q.

## **GLOBAL FOOTPRINT**



#### STUDY ABROAD EXPERIENCES 900+ INCOMING TO AUSTRALIA

4500+ OUTGOING TO PLACES AROUND THE WORLD

## OUR AUSTRALIAN







... is a data processing engine for Australian science. It empowers researchers to unlock impactful research discoveries within scientific data.

Integrative High Performance Computing

- usability by new HPC user communities over raw capacity;
- hardware suited to data processing;
- underpinning high performing wet, experimental and data-focused laboratories, with growing data processing needs;
- workflows that increase return on investment in instruments and experiments; and
- porosity and flexibility to serve specific requirements in growth usage areas, such as the life sciences, machine learning etc.



#### Professor Trevor Lithgow ARC Australian Laureate Fellow

Discovery of new protein transport machines in bacteria, understanding the assembly of protein transport machines, and dissecting the effects of antimicrobial peptides on anti-biotic resistant "super-bugs"



**FEI** Titan Krios

Nationally funded project to develop environments for Cryo analysis

Synchrotron

Data Management,

Integration between AS and MASSIVE M3

MX



MMI Lattice Light Sheet

Nationally funded project to capture and preprocess LLS data



MASSIVE M3 Structural refinement and analysis



Chamber details from the nanomachine that secretes the toxin that causes cholera. Research and data by Dr. Iain Hay (Lithgow Iab)

#### How do we partner with researchers... MASSIVE and Characterisation Virtual Laboratory





"Here is your CD of data..." "Your data is moving up to a data management system in the cloud where you have access to a range of tools and services to start your data analysis"





National Imaging



THE UNIVERSITY OF QUEENSLAND





#### How do we complement/replace workstations... Rich Online Workbenches



Atom Probe, **Structural Biology**, Bioinformatics, Cytometry, **Cryo-Electron Microscopy**, Neutron Beam Imaging, General Imaging Tool, Light Microscopy, General Scientific, X-ray









## What is Cryo Electron Microscopy? .. and some research context

## Pipeline in Biological Cryo-EM





## Computed Tomography + Photogrammetry



https://www.maximintegrated.com/en/app-notes/index.mvp/id/4682

http://www.cmu.edu/me/xctf/xrayct/index.html

## How many projections does a clinical CT use?



16 http://www.upstate.edu/radiology/education/rsna/ct/rec

## https://skfb.ly/TI89









3D reconstruction of the electron density of aE11 Fab' polyC9 complex

## What outcome are the scientists looking for?



Protein Data Bank in Europe Bringing Structure to Biology

Examples: hemoglobin, BRC/

#### PDBe > 5k12

Cryo-EM structure of glutamate dehydrogenase at 1.8 A resolution Source organism: *Bos taurus* 

**Primary publication:** 

Breaking Cryo-EM Resolution Barriers to Facilitate Drug Discovery.

Merk A, Bartesaghi A, Banerjee S, Falconieri V, Rao P, Davis MI, Pragani R, Boxer MB, Earl LA, Milne JL , Subramaniam S

Cell (2016) PMID: 27238019

Related structures: EMD-8194

Electron Microscopy 1.8Å resolution

Released: 08 Jun 2016



1.1

http://www.ebi.ac.uk/pdbe/entry/pdb/5k12

#### Structural Basis for Linezolid Binding Site Rearrangement in the Staphylococcus aureus Ribosome

#### Matthew J. Belousoff,<sup>a</sup> Zohar Eyal,<sup>b</sup> Mazdak Radjainia,<sup>c</sup> Tofayel Ahmed,<sup>d</sup> Rebecca S. Bamert,<sup>a</sup> Donna Matzov,<sup>b</sup> Anat Bashan,<sup>b</sup> Ella Zimmerman,<sup>b</sup> Satabdi Mishra,<sup>d</sup> David Cameron,<sup>a</sup> Hans Elmlund,<sup>c,e</sup> Anton Y. Peleg,<sup>a,f</sup> Shashi Bhushan,<sup>d,g</sup> Trevor Lithgow,<sup>a</sup> Ada Yonath<sup>b</sup>

Infection & Immunity Program, Biomedicine Discovery Institute & Department of Microbiology, Monash University, Clayton, Australia<sup>a</sup>; Department of Structural Biology, Weizmann Institute of Science, Rehovot, Israel<sup>b</sup>; The Clive and Vera Ramaciotti Centre for Structural Cryo-Electron Microscopy, Department of Biochemistry and Molecular Biology, Monash University, Victoria, Melbourne, Australia<sup>c</sup>; School of Biological Sciences, Nanyang Technological University, Singapore, Singapore<sup>d</sup>; Infection & Immunity Program, Biomedicine Discovery Institute & Department of Biochemistry and Molecular Biology, Monash University, Clayton, Australia<sup>e</sup>; Department of Infectious Diseases, Alfred Hospital, Prahran, Australia<sup>f</sup>; NTU Institute of Structural Biology, Nanyang Technological University, Singapore, Singapore<sup>g</sup>

**ABSTRACT** An unorthodox, surprising mechanism of resistance to the antibiotic linezolid was revealed by cryo-electron microscopy (cryo-EM) in the 70S ribosomes from a clinical isolate of *Staphylococcus aureus*. This high-resolution structural information demonstrated that a single amino acid deletion in ribosomal protein uL3 confers linezolid resistance despite being located 24 Å away from the linezolid binding pocket in the peptidyl-transferase center. The mutation induces a cascade of allosteric structural rearrangements of the rRNA that ultimately results in the alteration of the antibiotic binding site.

IMPORTANCE The growing burden on human health caused by various antibiotic

Received 27 March 2017 Accepted 17 April 2017 Published 9 May 2017

**Citation** Belousoff MJ, Eyal Z, Radjainia M, Ahmed T, Bamert RS, Matzov D, Bashan A, Zimmerman E, Mishra S, Cameron D, Elmlund H, Peleg AY, Bhushan S, Lithgow T, Yonath A. 2017. Structural basis for linezolid binding site rearrangement in the *Staphylococcus aureus* ribosome. mBio 8:e00395-17. https://doi.org/10 .1128/mBio.00395-17.

Disease resistance and drug targets

#### https://doi.org/10.1128/mBio.00395-17

ChemPubSoc Europe

Disease

resistance

and drug

targets

DOI: 10.1002/cmdc.201900042



#### cryoEM-Guided Development of Antibiotics for Drug-Resistant Bacteria

Matthew J. Belousoff,<sup>[a]</sup> Hari Venugopal,<sup>[b]</sup> Alexander Wright,<sup>[c]</sup> Samuel Seoner,<sup>[c]</sup> Isabella Stuart,<sup>[a]</sup> Chris Stubenrauch,<sup>[a]</sup> Rebecca S. Bamert,<sup>[a]</sup> David W. Lupton,<sup>\*[c]</sup> and Trevor Lithgow<sup>\*[a]</sup>

While the ribosome is a common target for antibiotics, challenges with crystallography can impede the development of new bioactives using structure-based drug design approaches. In this study we exploit common structural features present in linezolid-resistant forms of both methicillin-resistant *Staphylococcus aureus* (MRSA) and vancomycin-resistant *Enterococcus* (VRE) to redesign the antibiotic. Enabled by rapid and facile cryoEM structures, this process has identified (*S*)-2,2-dichloro-*N*-((3-(3-fluoro-4-morpholinophenyl)-2-oxooxazolidin-5-yl)methyl)acetamide (LZD-5) and (*S*)-2-chloro-*N*-((3-(3-fluoro-4-morpholinophenyl)) acetamide (LZD-6), which inhibit the ribosomal function and growth of linezolid-resistant MRSA and VRE. The strategy discussed highlights the potential for cryoEM to facilitate the development of novel bioactive materials.

dition to reaching resolution targets that allow unambiguous placement of small molecules, cryoEM is a solution-phase technique which may prove advantageous. Although cryoEMguided design of novel ligands is underdeveloped, the interactions of known drugs with their targets have been examined with the *Plasmodium* 20S proteasome,<sup>[6]</sup> the GPCR family,<sup>[7]</sup> pathogenic ribosomes,<sup>[8]</sup> receptor-bound insulin,<sup>[9]</sup> and important HIV viral entry proteins,<sup>[10]</sup> to name a few examples. In 2017, Scheres, Baum, and co-workers exploited cryoEM to elucidate the mode of action that mefloquine uses to inhibit the Plasmodium falciparum ribosome, and then used this insight to develop a next-generation molecule with enhanced antiparasitic activity.<sup>[11]</sup> As part of our studies into mechanisms of antibiotic resistance, we examined one of the escape routes that Staphylococcus aureus uses to develop resistance to the ribosomal-interfering antibiotic, linezolid.<sup>[8a]</sup> These studies, and those of others <sup>[12]</sup> show that disparate mutations result in a common

https://doi.org/10.1002/cmdc.201900042



# What is a time taken for analysis of a single dataset set?

## Paper Processing Time



## How much are GPUs used in preprocessing?



1	Pre	processing	<b>4</b>
	1.1	Getting organised	4
	1.2	Beam-induced motion correction	5
	1.3	CTF estimation	7
	1.4	Manual particle picking	9
	1.5	LoG-based auto-picking	1
	1.6	Particle extraction	3
	1.7	Making templates for auto-picking	4
	1.8	Selecting templates for auto-picking	8
	1.9	Auto-picking	9
		1.9.1 The shrink parameter	4
	1.10	Particle sorting	4

#### Important quote:

"The structure determination task is further **complicated** by the **lack of information** about the relative orientations of all particles and, in the case of **structural variability** in the sample, also their assignment to a structurally unique class."

#### 

#### A Bayesian View on Cryo-EM Structure Determination

#### Sjors H. W. Scheres

MRC Laboratory of Molecular Biology, Hills Road, Cambridge CB2 0QH, UK

Received 4 August 2011; received in revised form 27 October 2011; accepted 3 November 2011 Available online 12 November 2011

Edited by W. Baumeister

*Keywords:* cryo-electron microscopy; three-dimensional reconstruction; *maximum a posteriori* estimation Three-dimensional (3D) structure determination by single-particle analysis of cryo-electron microscopy (cryo-EM) images requires many parameters to be determined from extremely noisy data. This makes the method prone to overfitting, that is, when structures describe noise rather than signal, in particular near their resolution limit where noise levels are highest. Cryo-EM structures are typically filtered using *ad hoc* procedures to prevent overfitting, but the tuning of arbitrary parameters may lead to subjectivity in the results. I describe a Bayesian interpretation of cryo-EM structure determination, where smoothness in the reconstructed density is imposed through a Gaussian prior in the Fourier domain. The statistical framework dictates how data and prior knowledge should be combined, so that the optimal 3D linear filter is obtained without the need for arbitrariness and objective resolution estimates may be obtained. Application to experimental data indicates that the statistical approach yields more reliable structures than existing methods and is capable of detecting smaller classes in data sets that contain multiple different structures.

© 2011 Elsevier Ltd. Open access under CC BY license.

## What are the major analysis steps?

2D Classification 3D Classification Refinement

> Hint: All use GPUs and are done repeatedly.

## What is 2D Classification?

2D classification results are a way to visualize how your 2D images will cluster together into homogenous class averages without any model required.

#### **2D Classification**

The ML2D algorithm may be used to simultaneously align and classify single-particle images

An intrinsic characteristic of the ML approach is that it does not assign images to one particular class or orientation. Instead, images are compared with all references in all possible orientations and probability weights are calculated for each possibility.

(Maximum Likelihood)



Methods in Enzymology Volume 482, 2010, Pages 321-341

M	ethe V M	als i	in NCN	
	dana 	er Hi L Part		

Chapter Twelve - Methods for Three-Dimensional Reconstruction of Heterogeneous Assemblies

Elena V. Orlova, Helen R. Saibil

**E** Show more

https://doi.org/10.1016/S0076-6879(10)82013-0

Get rights and content

https://doi.org/10.1016/S0076-6879(10)82013-0

#### Particle picking for next step of classification in 2D



Source: https://www.e-sciencecentral.org/articles/Table.php?xn=am/am-48-001&id=

### **Classification of selected particles in 2D**



## What is 3D Classification?

"The premise of 3D classification is to use an initial starting model to sort your particles into 3 or more groups. By using the 3D model, Relion uses maximum likelihood methods to create homogenous classes of particles that will belong to a given group."



# What is a time taken for analysis of a single dataset set?

#### What is the problem? Antibiotics at crisis point

Slow time from lead compound to clinic Increasing prevalence of resistant bacteria in the clinic

UK AMR review: "Deaths by infection will be larger than both cancer and diabetes combined by 2050....."

A big problem is Methicillin-resistant Staphylococcus aureus (MRSA) and Vancomycin-resistant Enterococci (VRE)...

nulles.										
Pubymed.gov	PubMed 🗘									
US National Library of Medicine National Institutes of Health		Advanced								
Format: Abstract -		Ser	nd to 👻							
Eur Rev Med Pharmacol Sci. 20	17 Oct;21(17):3974-3979	9.								
Linezolid vs. vancomycin in treatment of methicillin-resistant staphylococcus aureus infections:										
a meta-analysis.										
Li J <sup>1</sup> , Zhao QH, Huang KC, Li ZQ, Zhang LY, Qin DY, Pan F, Huang WX.										
Author information										
Abstract										
OBJECTIVE: This study aims to explore the treatment of methicillin-resistant staphylococcus aureus (MRSA) infection by using meta-analysis										
method.										
MATERIALS AND METHO	DS: Pubmed/Medlin	ne, ScienceDirect, CNKI and Wanfang database were comprehensively searched to obtain	the							
randomized controlled trials (RCTs) on linezolid and vancomycin in the treatment of MRSA infections. We extracted features and information										
of included studies and selected appropriate effect models based on the heterogeneity test results. The funnel plot was used to analyze										
publication bias.										
RESULTS: A total of sever	1 RCTs including 537	76 cases met the inclusion criteria. Meta-analysis showed that the clinical cure rate of linez	zolid							
group was higher than that	t of vancomycin grou	up after treatment (OR = 1.85; 95% CI: 1.33-2.59, p<0.001) and follow-up (OR = 1.49; 95%	6 CI:							
1.17-1.91, p=0.001). In the microbiologically evaluable patients, end of therapy (EOT) MRSA clearance rate, and test of cure (TOC) MRSA										
clearance rate of linezolid	were superior to tho	ose of vancomycin.								

**CONCLUSIONS:** Based on the combined analysis of randomized controlled trials, the efficacy of linezolid should be better than that of vancomycin in the treatment of infections caused by MRSA, but conclusions still need to be further validated by more well-designed RCTs of a large sample.


## Patient to Microscope "Pragmatic Approach"

- Rational redesign of existing drug pharmaphacores
  - Linezolid
  - Virginiamycin (Synercid)
    - More tractable chemical synthesis of fully synthetic derivatives
  - Thiostrepton
    - Semi-synthetic modifications to help with drug solubility



Electron source (Field Emission Gun)



2D Particle Classes



#### TEM Micrograph



- CTF Correction of micrographs
- Particle extraction
- Classifcation of Particles
- -~300 500k particles collected





- Manual model building (coot)
- Refinement (Phenix)

### 2017

- ~1-4TB raw data set/sample ~2000-5000 files
- Pipeline analysis with internal & external tools
- Require large memory gpu > 8GB
- Require large system memory > 64GB
- Require cpu cores 200 400
- Parallel file reads and writes

## 2019

How long

does an

analysis take

(and why?)?

- ~1-10TB raw data set/sample ~2000-10000 files
- Require large memory gpu > 16GB
- Require large system memory > 18GB
- Require cpu cores 30 130
- Parallel file reads and writes and high ops local cache

Task	Submitted?	GPU?	Nodes	Time	
Import	No			< 1 min	
Motion Correction	Yes	Yes	3	20 min	
CTF estimation	Yes	No	1	20 min	
Manual Picking	No			?	
Autopicking	Yes	Yes	2	40 min	
Particle Extraction	Yes	No	1	10 min	
2D Classification	Yes	Yes	2	10 min/iteration	
<b>3D</b> Classification	Yes	Yes	1	10 min/iteration	
3D Refine	Yes	Yes	2	5-10 min/iteration	
Movie Refine	Yes	No	1	1 hour	
Particle Polishing	Yes	No	1	1-2 hours	
Mask Creation	No			5-30 min	
Postprocessing	No			<1 min	

How long do each of the steps take?

2,500 images

150,000 particles

260 pixels

## **Malaria Parasitic Protein**

Protein used for metabolic recycling ~200 kDa *150 Å diameter* Resolution = 2.6 Å Particles Used in Reconstuction = 32 k

Unpublished Data

How long does an analysis take (and why?)?





## Does GPU architecture affect performance?

~ 2 x per generation ( Kepler -> Volta)



## Does job layout matter between GPU architectures?

Volta peaks at lower MPI ranks, leaves more GPU Ram available



## How does GPU Ram affect scientific outcomes?

K3 cameras produce 5k images



## Workstation GPUs vs HPC GPUs?

~30% Faster, even faster for higher precision



## **Methicillin Resistant Ribosome**

Nanomachine used for protein synthesis In complex with new antibiotic - Tedizolid ~3.8 MDa *310 Å diameter* Resolution = 3.1 Å Particles Used in Reconstuction = 77 k

Unpublished Data

Job\_number 40 35 **GPU** acceleration 30 for 11 of 30 **b** 25 **c** 20 **c** 15 10 14 DAYS 5 2018-11-27 2018-11-29 2018-12-01 2018-12-03 2018-12-05 2018-12-07 2018-12-09 2018-12-11 **Folder Create Time** 

How long does an analysis take (and why?)? How long does an analysis take (and why?)?

















## How fast can we process this data?

... hardware choices, do they matter?

## **Hardware Configurations**

## Workstation



## 8 Cores 64GB Ram 4 x GTX1070

SSD + NAS

Cloud (HPC)



Dual Socket 36 Cores 384GB Ram 3 x V100



NVIDIA DGX-1 DEEP LEARNING SYSTEM

Dual Socket 40 Cores 512GB Ram 8 x V100

SSD + Lustre





#### 2D Classification Step Runtime for Different Hardware and Optimisations

Local scratch on **OPTIONS** MPI result combine off **OPTIMISATION** Parallel disk access off Pre-read particles off User config 0.00 10.00 70.00 20.00 30.00 40.00 50.00 60.00 80.00 **JOB RUNTIME (MIN)** ■ Hardware Type 1xm3g Hardware Type DGX1-V

Can we process faster? (do options matter?)

#### 2D Classification Step Runtime for Different Hardware and Optimisations

Can we process faster? (do options matter?)



#### Cost and Performance Ratios for Workstation VS HPC



What is the ratio of cost to performance?



# What retail options for compute exist?

Workstations + Software Support

68

How fast are the Relion developers machines?

- [lancew@m3-login1 job008]\$ pwd
- /scratch/br76/relion3tut/relion30\_tutorial\_precalculated\_results/Class2D/job008
- [lancew@m3-login1 job008]\$ ls --full-time
- total 2884
- -rw-r--r-- 1 lancew br76 2777 2018-06-07 23:32:58.000000000 +1000 default\_pipeline.star
- -rw-r--r-- 1 lancew br76 840 2018-06-07 23:32:58.000000000 +1000 job\_pipeline.star
- -rw-r--r-- 1 lancew br76 448 2018-06-07 23:32:58.000000000 +1000 note.txt
- -rw-r--r-- 1 lancew br76 0 2018-06-07 23:32:57.000000000 +1000 run.err
- -rw-r--r-- 1 lancew br76 820224 2018-06-07 23:32:58.000000000 +1000 run\_it025\_classes.mrcs
- -rw-r--r-- 1 lancew br76 740401 2018-06-07 23:32:57.000000000 +1000 run\_it025\_data.star
- -rw-r--r-- 1 lancew br76 223301 2018-06-07 23:32:57.000000000 +1000 run\_it025\_model.star
- -rw-r--r-- 1 lancew br76 5916 2018-06-07 23:32:58.000000000 +1000 run\_it025\_optimiser.star
- -rw-r--r-- 1 lancew br76 458 2018-06-07 23:32:56.000000000 +1000 run\_it025\_sampling.star
- -rw-r--r-- 1 lancew br76 1234 2018-06-07 23:32:57.000000000 +1000 run.job
- -rw-r--r-- 1 lancew br76 307687 2018-06-07 23:32:58.000000000 +1000 run.out
- -rw-r--r-- 1 lancew br76 820224 2018-06-07 23:32:58.000000000 +1000 run\_unmasked\_classes.mrcs

#### **DO MORE SCIENCE! PLUG AND PLAY** Have peace of mind, focus on what matters Exxact systems are fully Suggested R most, knowing your system is backed by a 3 turnkey, built to perform year warranty and support. right out of the box so you avoid the drudgery of configuration and setup. Why do these **Entry-Level** Mid-Range **High-End** systems look **Tensor Workstation Tensor Workstation Tensor Server** attractive to 1x Intel Core i7-7820X 2x Intel Xeon Scalable (Silver) 2x Intel Xeon Scalable (Gold) researchers? 4x NVIDIA RTX 2080 Ti, RTX 2080 or 4x NVIDIA RTX 2080 Ti, RTX 2080 or Up to 8x NVIDIA Tesla GPUs GTX 1080 Ti GTX 1080 Ti 12x 32GB Memory 🚟 8x 16GB Memory I2x 16GB Memory 1x 512GB M.2 SSD (OS) 1x 256GB SSD (OS) 2x 240GB SSD (OS) 2x 2TB HDD (Data) 2x 2TB SSD (Data) 2x 2TB HDD (Data) STARTING AT \$9,599 STARTING AT \$12,999 STARTING AT \*\$22,999

**START YOUR ORDER** 

Source: https://www.exxactcorp.com/Relion-for-Cryo-EM-Solutions

**START YOUR ORDER** 

**START YOUR ORDER** 

#### PRECONFIGURED

"Relion 3.0 beta and v2.1 With example job submission scripts, benchmarks, a fully validated test suite, and the latest software patches for quick implementation."

> Why do these systems look attractive to researchers?



Source: https://www.exxactcorp.com/Relion-for-Cryo-EM-Solutions

Why do these systems look attractive to researchers?

#### 4. Turn-key Systems.

All of our systems are delivered ready to use inmediately. We setup the Workstation parameters such as Networking, Users creation, Backup automated sytem,

Users configuration.

Networking, hostname, DNS's, etc....

Services.

Data migration

Virtual Machines Installations

Remote access configuration.

Automated Data Backup

Source: http://www.linuxvixion.com/wp-content/uploads/2016/04/linuxvixion-molecular-dynamics-2016-english.pdf

## **Amazon EC2 P3 Instance Product Details**

Instance Size	GPUs - Tesla V100	GPU Peer to Peer	GPU Memory (GB)	vCPUs	Memory (GB)	Network Bandwidth	EBS Bandwidth	On-Demand Price/hr*	1-yr Reserved Instance Effective Hourly*	3-yr Reserved Instance Effective Hourly*
p3.2xlarge	1	N/A	16	8	61	Up to 10 Gbps	1.5 Gbps	\$3.06	\$1.99	\$1.05
p3.8xlarge	4	NVLink	64	32	244	10 Gbps	7 Gbps		\$7.96	\$4.19
p3.16xlarge	8	NVLink	128	64	488	25 Gbps	14 Gbps	\$25/hr	\$15.91	\$8.39
p3dn.24xlarge	8	NVLink	256	96	768	100 Gbps	14 Gbps		\$18.30	\$9.64

\* - Prices shown are for Linux/Unix in the US East (Northern Virginia) AWS Region and rounded to the nearest cent. For full pricing details, see the Amazon EC2 pricing page.

Customers can purchase P3 instances as On-Demand Instances, Reserved Instances, Spot Instances, and Dedicated Hosts.

## What are the cloud options? ...and cost?

Source: http://labphoto.tumblr.com/post/105112424006/cyanuric-triazide-or-2-4-6-triazido-1-3-5-triazine


# What has changed in 2 years?

Mostly hardware .... And a little ML

75

# **Options for Computation (2017)**

Workstation

Pro

Full user control

Con

Limited by single machine

Cloud Pro Scales easily Con Cost, complexity, data movement

HPC Pro Huge resources Con Tightly controlled, shared

# **Options for Computation (2019)**

# Workstation

Pro

Full user control

Con

Limited by single machine

Cloud

Scales easily

Con

Cost, complexity, data movement

HPC Pro Huge resources, purpose designed Con Tightly controlled, shared

Containers



#### SYSTEM SPECIFICATIONS

GPUs	8X Tesla V100
Performance (Mixed Precision)	1 petaFLOPS
GPU Memory	256 GB total system
CPU	Dual 20-Core Intel Xeon E5-2698 v4 2.2 GHz
NVIDIA CUDA <sup>®</sup> Cores	40,960
NVIDIA Tensor Cores (on V100 based systems)	5,120
Power Requirements	3,500 W
System Memory	512 GB 2,133 MHz DDR4 RDIMM
Storage	4X 1.92 TB SSD RAID 0
Network	Dual 10 GbE, 4 IB EDR
Operating System	Canonical Ubuntu, Red Hat Enterprise Linux
System Weight	134 lbs
System Dimensions	866 D x 444 W x 131 H (mm)
Packing Dimensions	1,180 D x 730 W x 284 H (mm)
Operating Temperature Range	5–35 °C



GPUs	8X Tesla V100	8X Tesla P100			
TFLOPS (GPU FP16)	960	170			
GPU Memory	128 GB to	tal system			
CPU	Dual 20-Core Intel Xeon E5-2698 v4 2.2 GHz				
NVIDIA CUDA® Cores	40,960	28,672			
NVIDIA Tensor Cores (on V100 based systems)	5,120	N/A			
Maximum Power Requirements	3,200 W				
System Memory	512 GB 2,133 MHz DDR4 LRDIMM				
Storage	4X 1.92 TB	SSD RAID 0			
Network	Dual 10 Gb	E, 4 IB EDR			
Software	Ubuntu Lin See Software S	ux Host OS tack for Details			
System Weight	134	lbs			
System Dimensions	866 D x 444 W	x 131 H (mm)			
Packing Dimensions	1,180 D x 730 V	V x 284 H (mm)			
Operating Temperature Range	10–3	85 °C			



GPUs	8x Tesla P100
TFLOPS (GPU FP16 / CPU FP32)	170/3
GPU Memory	16 GB per GPU
CPU	Dual 20-core Intel® Xeon® E5-2698 v4 2.2 GHz
NVIDIA CUDA® Cores	28672
System Memory	512 GB 2133 MHz DDR4
Storage	4x 1.92 TB SSD RAID 0
Network	Dual 10 GbE, 4 IB EDR
Software	Ubuntu Server Linux OS DGX-1 Recommended GPU Driver
System Weight	134 lbs
System Dimensions	866 D x 444 W x 131 H (mm)
Packing Dimensions	1180 D x 730 W x 284 H (mm)
Maximum Power Requirements	3200W
Operating Temperature	10 - 30°C

## P100 -> V100

## Relion - 3D Classification Step (2017)



# Relion - 3D Classification Step (2019)



### Relion Class 3D - MPI Task Effects (2017)



### Relion Class 3D - MPI Task Effects (2019)



#### **Iteration Processing Time VS Number of DGX-1 Servers**



45



# What has stayed the same in 2 years?

Mostly software .... Still a struggle

#### openmpi-bin\_2.1.1-8\_amd64.deb

#### Description

openmpi-bin - high performance message passing library -- binaries

Property	Value
Distribution	Ubuntu 18.04 LTS (Bionic Beaver)
Repository	Ubuntu Universe amd64
Package name	openmpi-bin
Package version	2.1.1
Package release	8
Package architecture	amd64
Package type	deb
Installed size	444 B
Download size	86.12 KB
Official Mirror	archive.ubuntu.com



bforsbe commented 5 hours ago

Contributor + 😐 🚥

openmpi 2.0.2 is my go-to version, for future reference.

#### Software Dependencies



# How to accelerate a solution with HPC?

Using a real example (K2)

j	The number of parallel threads to run on each CPU. <b>We often use</b> <b>4-6.</b>
 dont_combine_weights_ via_disc	By default large messages are passed between MPI processes through reading and writing of large files on the computer disk. By giving this option, the messages will be passed through the network instead. <b>We often use this option.</b>
no_parallel_disc_io	By default, all MPI slaves read their own particles (from disk or into RAM). Use this option to have the master read all particles, and then send them all through the network. We do not often use this option.
preread_images	By default, all particles are read from the computer disk in every iteration. Using this option, they are all read into RAM once, at the very beginning of the job instead. <b>We often use this option if</b> <b>the machine has enough RAM</b> (more than N*boxsize*boxsize*4 bytes) to store all N particles.
scratch_dir	By default, particles are read every iteration from the location specified in the input STAR file. By using this option, all particles are copied to a scratch disk, from where they will be read (every iteration) instead. We often use this option if we don't have enough RAM to read in all the particles, but we have large enough fast SSD scratch disk(s) (e.g. mounted as /tmp).

What are the optimisation options?

What are the optimisation options?







How are the CPUs used? ... How many are needed?



How are the CPUs used? ... How many are needed?



How are the CPUs used? ... How many are needed?



### How efficient is the software... GPU Performance Measurements

Utilization rates report how busy each GPU is over time, and can be used to determine how much an application is using the GPUs in the system.

GPU Percent of time over the past sample period during which one or more kernels was executing on the GPU.

Memory Percent of time over the past sample period during which global (device) memory was being read or written.

PCIe Rx and Tx Throughput in MB/s

MONASH

<sup>erSity</sup> https://developer.download.nvidia.com/compute/DCGM/docs/nvidia-smi-367.38.pdf

Latest generation GPU hardware enables very detailed performance metrics.















How much power is consumed ? ... global warming?



How cooling is needed ? ... global warming?

Relion3 Tutorial Dataset



# Where are the opportunities for reducing "Time to Science"?

### Where is development heading... Relion 3 – CPU Vs GPU?

However, the GPU acceleration is only available for cards from a single vendor, and it cannot use many of the largescale computational resources available in existing centres, local clusters, or even researchers' laptops. In addition, the memory available on typical GPUs limits the box sizes that can be used, which could turn into a severe bottleneck for large particles. For relion-3, we have developed a new general code path where CPU algorithms have been rewritten to mirror the GPU acceleration, which provides dual execution branches of the code that work very efficiently both on accelerators as well as the single-instruction, multiple-data vector units present on traditional CPUs.

Limitations of memory

Increased performance from new CPUs



# Opportunities for speed up

Cryosparc
 Cryolo
 Preprocessing tools

#### 之前 cryoSPARC

Features Documentation Blog Download

# Unlock the potential of cryo-EM with cryoSPARC<sup>™</sup>

CryoSPARC is the state-of-the-art platform used globally for obtaining 3D structural information from single particle cryo-EM data.

The cryoSPARC platform enables automated, high quality and highthroughput structure discovery of proteins, viruses and molecular complexes for research and drug discovery.

**Download cryoSPARC™** 

Version 2.5.0 Released: January 8, 2019



Source: https://cryosparc.com/

# Opportunities for speed up

Cryosparc
 Cryolo
 Preprocessing tools



Source: http://sphire.mpg.de/wiki/doku.php?id=downloads:cryolo\_1&redirect=1

# Known issues

- Issue 0: Training on multiple GPUs sometimes lead to worse performance (higher loss). We currently recommend to train on single gpus.
- Issue 1: crYOLO sometimes not exit properly after training finished. Has to be killed manually.
- Issue 2: If you use automatic filtering with .tif files, you get an error like "OSError: cannot identify image file 'filtered\_folder/another\_folder/my\_image.tif'". It will be fixed in the next release.
- Issue 3: (Boxmanager) The visualization only shows the first filament when loading eman1 helical box files (start end coordinates). Will be fixed in the next release.
- Issue 4: The filament mode will crash if crYOLO cannot identify a single particle in the image. Will be fixed in 1.2.2
- Issue 5: If movies were aligned with cisTEM and picked with crYOLO, the box position are vertically flipped. Will be fixed in 1.2.2
- Issue 6: crYOLO does overwrite the environmental variable "CUDA\_VISIBLE\_DEVICES" with 0 if no gpu is specified by the -g
  parameter. This leads to the behavior that crYOLO ignores previous settings in CUDA\_VISIBLE\_DEVICES. Will be fixed in 1.2.2
- Issue 7: On K3 images crYOLO seems to add a offset toward the longer axis of the input image.
- Issue 8: There is a logical error in filament tracing, which sometimes connects two parallel filaments.
- Issue 9: Some people report an error when running cryolo prediction/training: "ImportError: numpy.core.multiarray failed to import". It will be fixed in 1.2.3.
- Issue 10: On machines with many cores (e.g 64) an error during filtering might pop up: "[ERROR:0] 53: Can't spawn new thread"
- Issue 11: If the -g parameter is not provided, crYOLO will use the memory of all GPUs. Will be fixed in 1.2.3.
- Issue 12: The LineEnhancer dependenceny of crYOLO is still dependent from opency. Workaround: In the crYOLO environment: conda install opency
- Issue 13: After picking it can happen that some of the boxes are not fully immersed in the image.

Opportunities for speed up

 Cryosparc
 Cryolo
 Preprocessing tools

Characterisation-Virtual-Labo	O Unwatch ▼	1 🕇 Star 1	% Fork 0						
<>Code ① Issues 0     ी Pull r	requests 0 🕅 Projects	0 🗉 Wiki	Insights	🌣 Settings					
Tools for optimising preprocessing of cryoem micrographs           Edit           Manage topics									
🕝 54 commits	្រំ <b>1</b> branch	🛇 <b>1</b> release	<u>11</u> :	2 contributors	مِٹِه GP	L-3.0			
Branch: master - New pull request			Create new fi	le Upload files	Find file Clone	or download 🔻			
vzej8y Updated user guide and install documentation Latest commit ccc8748 on 3 Jan									
images Updated user guide and install documentation 22						2 months ago			
.gitignore	Documentation Update	Documentation Updated				3 months ago			
Cryo-EM.ipynb	Workflow debug set to false					2 months ago			
	Initial commit 5 mont								
Minutes.txt	Updated user guide and install documentation 2 mo					2 months ago			
README.md	Altered text	Altered text				3 months ago			

https://github.com/Characterisation-Virtual-Laboratory/Cryo-EM-Processing-Tool

#### Cryo-EM Processing Tool

Motion Correction	Contrast Transfer Auto Particle Picking			
Basic	Advanced			
Job no				
Input	mrc file, folder containing mrc files or .star file			
Output	NBMotionCorr/			
Pixel size (A)			0.50	
Patch		Patch list:	Use ";" as separator. e.g. 3 3; 5 5;	✓ Add jobs
B-Facto	· 150	bFactor list:	Use ";" as separator. e.g. 100; 235; 350	✔ Add jobs
Voltage (kV)	: 300			
Gair	path for mrc file that stores the gain reference			
GPU usage	indicate the GPUs to use			
Frame dose (e/A2)			1.0	
🗸 Run				
Jobs: Job#   inMrc	outMrc   pixelSize   patch   bFactor   voltage   gainFile   gpu   inTiff   fullSu	um   defectFile   processi	ng   iteration   tolerance   stack   binningFactor   initDose   frameDose	throw
Add	Delete Select Update	Run All	Progress:	
	https://github.com/Charact	erisation-Virt	ual-Laboratory/Cryo-EMI-Processing-	
	Tool			

Ω

CTF Review											
Start review	Erro	rs									
Project directory:	/Users/j	/Users/jvanschyndel/Documents/Relion/betagal/				Folder:	NBCtfFind/				
Available jobs:	● ctf1 ● ctf2	ctf1 ctf2				Total micrographs	ctf1: ctf2:	30 180			
Filters											
		•									
Defocus	U from:	0					to:	99999999			
Defocus	s v from:	0					to:	99999999			
Defocus ang	gle from:	0					to:	99999999			
Volta	ge from:	0					to:	99999999			
Spherical aberrat	tion fro	0					to:	99999999			
Amplitude contra	ast from:	0					to:	99999999			
Magnificati	on from:	0					to:	99999999			
Detector pix si	ze from:	0					to:	99999999			
Ctf figure of me	rit from:	0					to:	99999999			
Ctf max resolution	on from:	0					to:	99999999			
Reset		max resolu Appl	t ition y								
✓ Save											
				) )		)			$\sum_{i=1}^{n}$		).
✓ Data	✓ Rai	w mrc	🗸 Data	🗸 Raw mrc	✔ Data	✓ Raw mrc		🖌 Data	🛩 Raw mrc	🛩 Data	🗸 Raw mrc
Defocus V: Defocus V: Defocus Y: Yoltage: Spherical Aberration: Amplitude Contrast: Magnification: Detector Pixel Size: Figure of Merit: Ctf Max Resolution:	5007.8 5074.23 13.47 300.0 : 2.7 0.1 104477. 14.0 0.05745 19.542	6	https://git	wh com/C	aracterisat	ion-Virtual		noratory		ocessing	
			https://giti		alacterisat		Lar	Joratory,	CIYU-LIVI-FI	occasing-	
			Tool								
#### **Cryo-EM Processing Tool**

Workflow										
			_							
Mode:	Single	Workflow								
Project Directory:	/Users/jvanschyndel/Documents/Relion/betagal/									
🗸 Load jobs	✓ Save j	jobs Erron								
🗸 Run all jobs	Progress:									
Auto run:	O Disable	St	tart auto run	Stop auto run	Pause (mins):	5		Running (secs):	0	
	Enable									
Display graphs:	Disable									
	O Enable									

https://github.com/Characterisation-Virtual-Laboratory/Cryo-EM-Processing-Tool

2



### Next Generation: 4D+ Volumetric Imaging Lattice light sheet microscopy



Inset images show the key stages as captured by lattice lightsheet microscopy, structured illumination microscopy, and correlative light and electron microscopy.



#### Mitochondrial DNA (green) escaping mitrochondria (red) during cell death.

Dr Kate McArthur (Monash BDI) and Dr Lachlan Whitehead (WEHI) and The Advanced Imaging Centre at Janelia Research Campus.





Credit: Steve Morton

## Acknowledgments

Dr Matt Belousof Hari Venogopal Jafar Lie Jay Van Schyndel

MASSIVE Partners ARDC



# THANK YOU

Final messages:

- ✓ HPC is the best option, if you have interactive queues.
- $\checkmark$  Cloud is good for one off analysis.
- Workstations make great training/learning platforms

For future questions: lance.wilson@monash.edu