HIGH-PERFORMANCE BIOLOGICAL COMPUTING

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Issues in Computational Genomics

A view from the trenches

C. Victor Jongeneel
Director, HPCBio group
Institute for Genomic Biology and
National Center for Supercomputing Applications

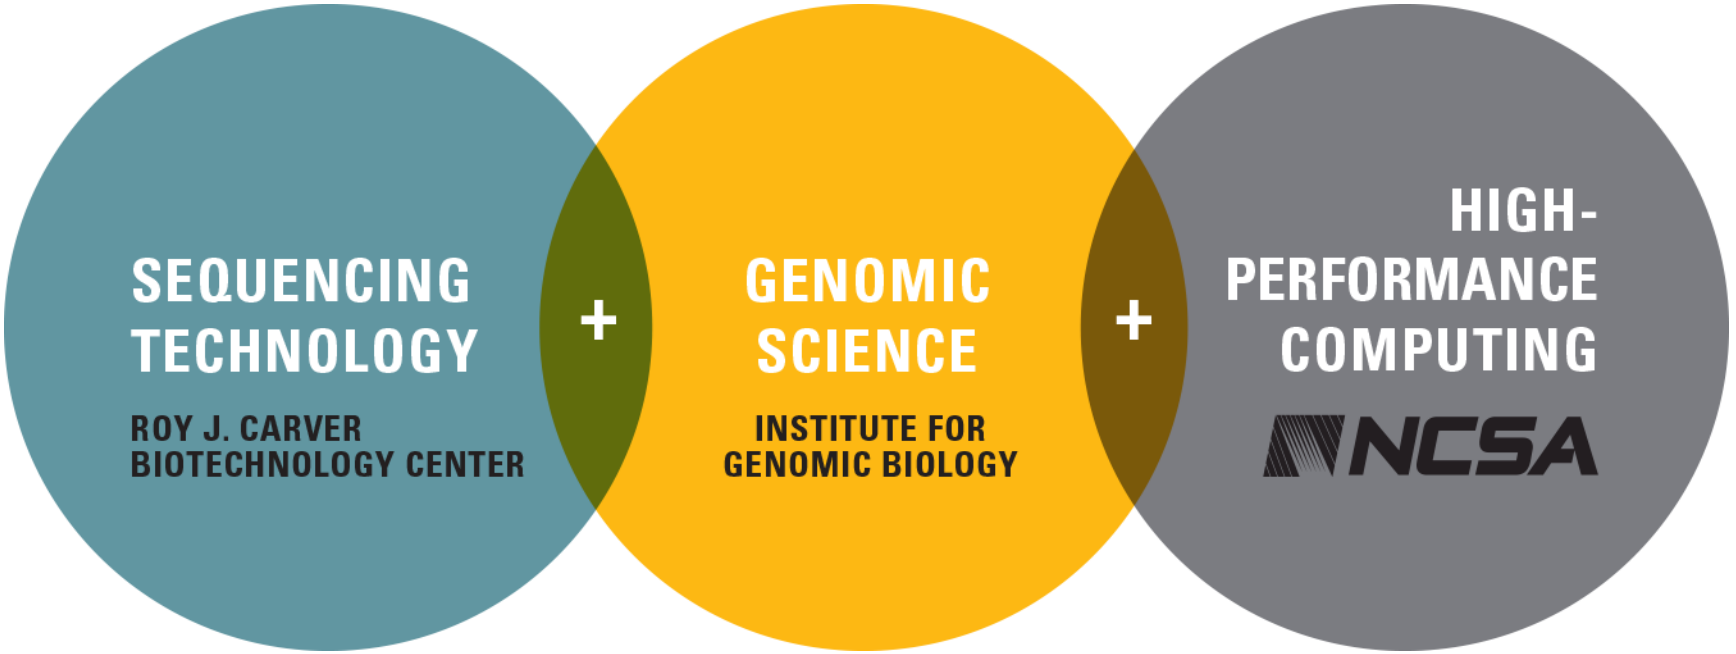
SEQUENCING TECHNOLOGY
ROY J. CARVER BIOTECHNOLOGY CENTER
+
GENOMIC SCIENCE
INSTITUTE FOR GENOMIC BIOLOGY
+
HIGH-PERFORMANCE COMPUTING
NCSA

Missions of HPCBio

- Infrastructure
  - » Robust computational infrastructure
  - » Software and databases
- Support
  - » Experimental design and consulting
  - » Analysis of high-throughput data
- Training
  - » Organization of workshops and short courses
- Applied R&D
  - » Development of innovative solutions to problems in computational genomics, driven by customer projects

IGB HPC resources

Default
- Where most jobs are run
- 10 – Dell R410s
- 80 Cores, 240GB RAM

Classroom
- Supports teaching lab
- 24 – Dell 2950 III
- 192 cores, 384 GB RAM

Large-Memory
- Genome assemblies, other applications that need more RAM
- 1 – Dell R910
- 24 cores, 1TB RAM

Blacklight
- Large jobs, genome and transcriptome assemblies
- 4 nodes – SGI UV 1000
- 1568 Cores, 8 TB RAM

Storing and archiving data

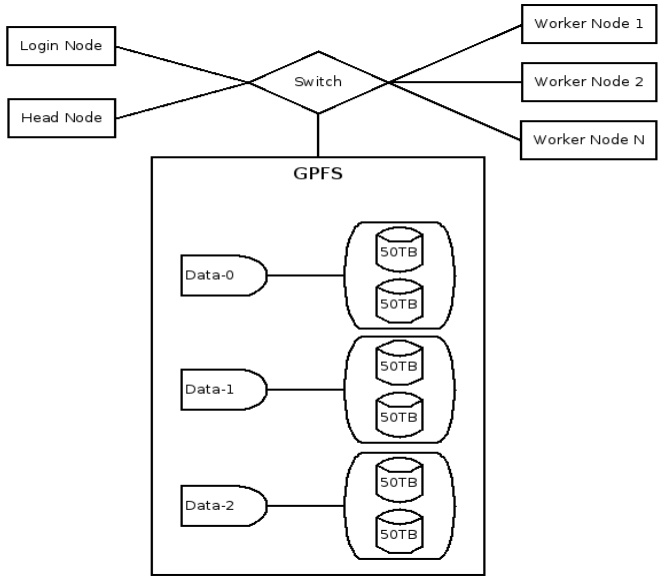

Automated “drag & drop” archiving of petabytes of data onto high-density tapes

Recent grant from Roy J. Carver Charitable Trust

General Parallel File System (GPFS) allows access from all cluster nodes with very high throughput

iForge

Resource for the NCSA Private Sector Program

Three components:

- Large (2000 cores) blade center ☒
- Fat nodes (32 cores, 256 GB RAM)
- Large L3 cache nodes

Community requirements

Mid-size DNA sequencing facility serving a large research campus and many off-campus customers

- Two 454 Titanium, three HiSeq 2000, one MiSeq, all running essentially 24/7
- Extremely diverse user community

Alliance with Mayo Clinic

- Planning to sequence thousands of genomes (exomes or whole) per year
- Do not have the computational resources to process all their data locally

Some analyses performed by HPCBio

- Genome and transcriptome *de novo* assembly, from bacteria to complex eukaryotes
- Genome and transcriptome annotation (structure and function of genes)
- Analysis of RNA-seq and ChIP-seq data
- Analysis of metagenomes (WGS) and of community structures (amplicons)
- Discovery and calling of genetic variants in individuals and in populations
- Statistical analysis of gene expression data (RNA-seq and microarrays)
- Custom scripting and programming

A typical data-intensive workflow

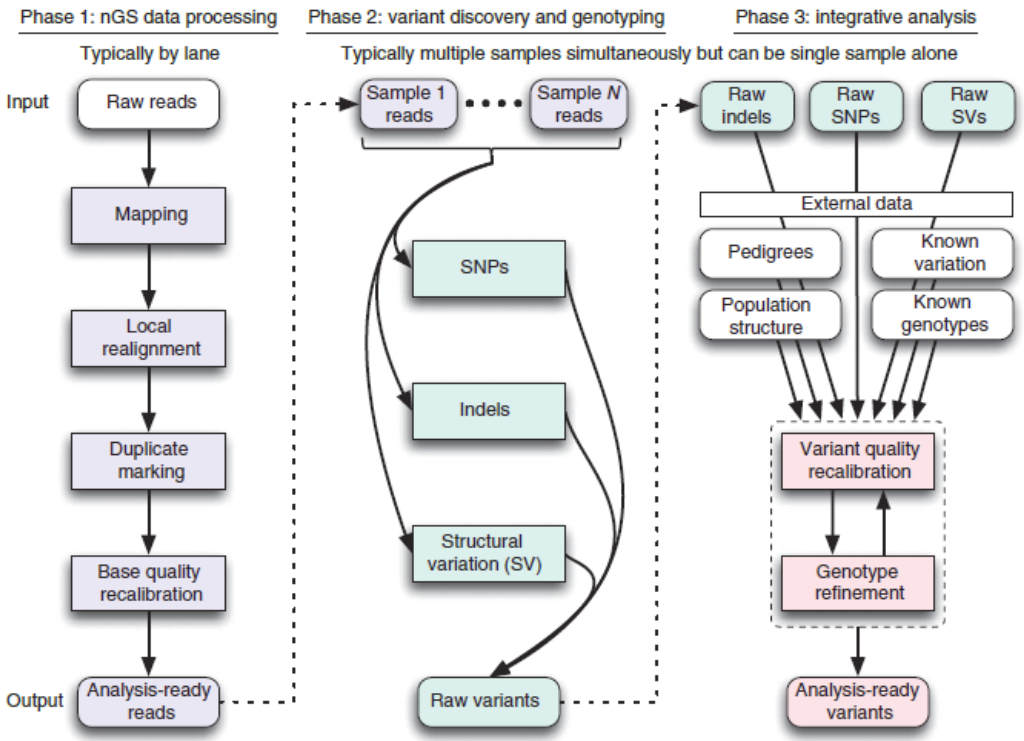

From DePristo et al, Nature Genetics, May 2011

Variant calling and analysis workflow used for the 1000 Genomes project

Input: up to 1.5×10^9 Illumina reads (>500 GB) plus a reference human genome

Outputs: many files – BAM, VCF, plain text

△ One may want to run many of these workflows simultaneously !

Data-level parallelization

Typical strategy when there are no dependencies between subsets of the data

For variant calling, start with producing chunks of sequence reads and map each chunk to the reference genome, then merge all alignments and split again per chromosome before doing recalibration and variant calling

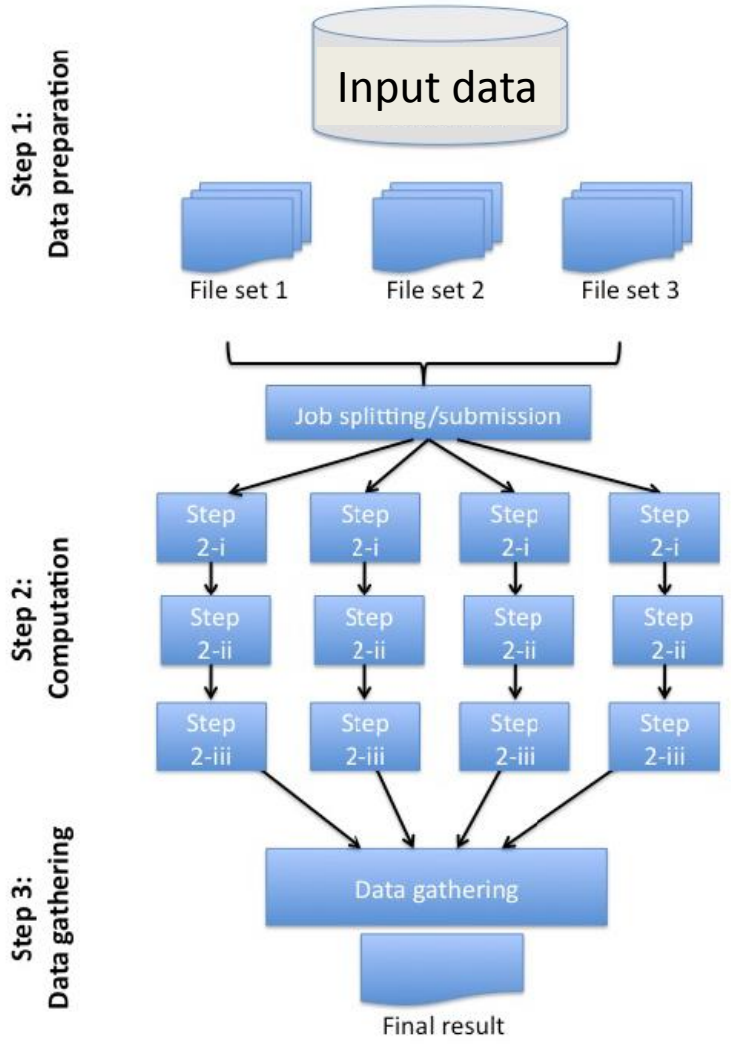

A real life challenge

Mayo Clinic had set of 225 human exome sequence datasets that required immediate analysis

They had written a workflow wrapper with no data level parallelization and no error management

Initial tests showed that job failures were in fact quite common and caused catastrophic failures of the entire workflow

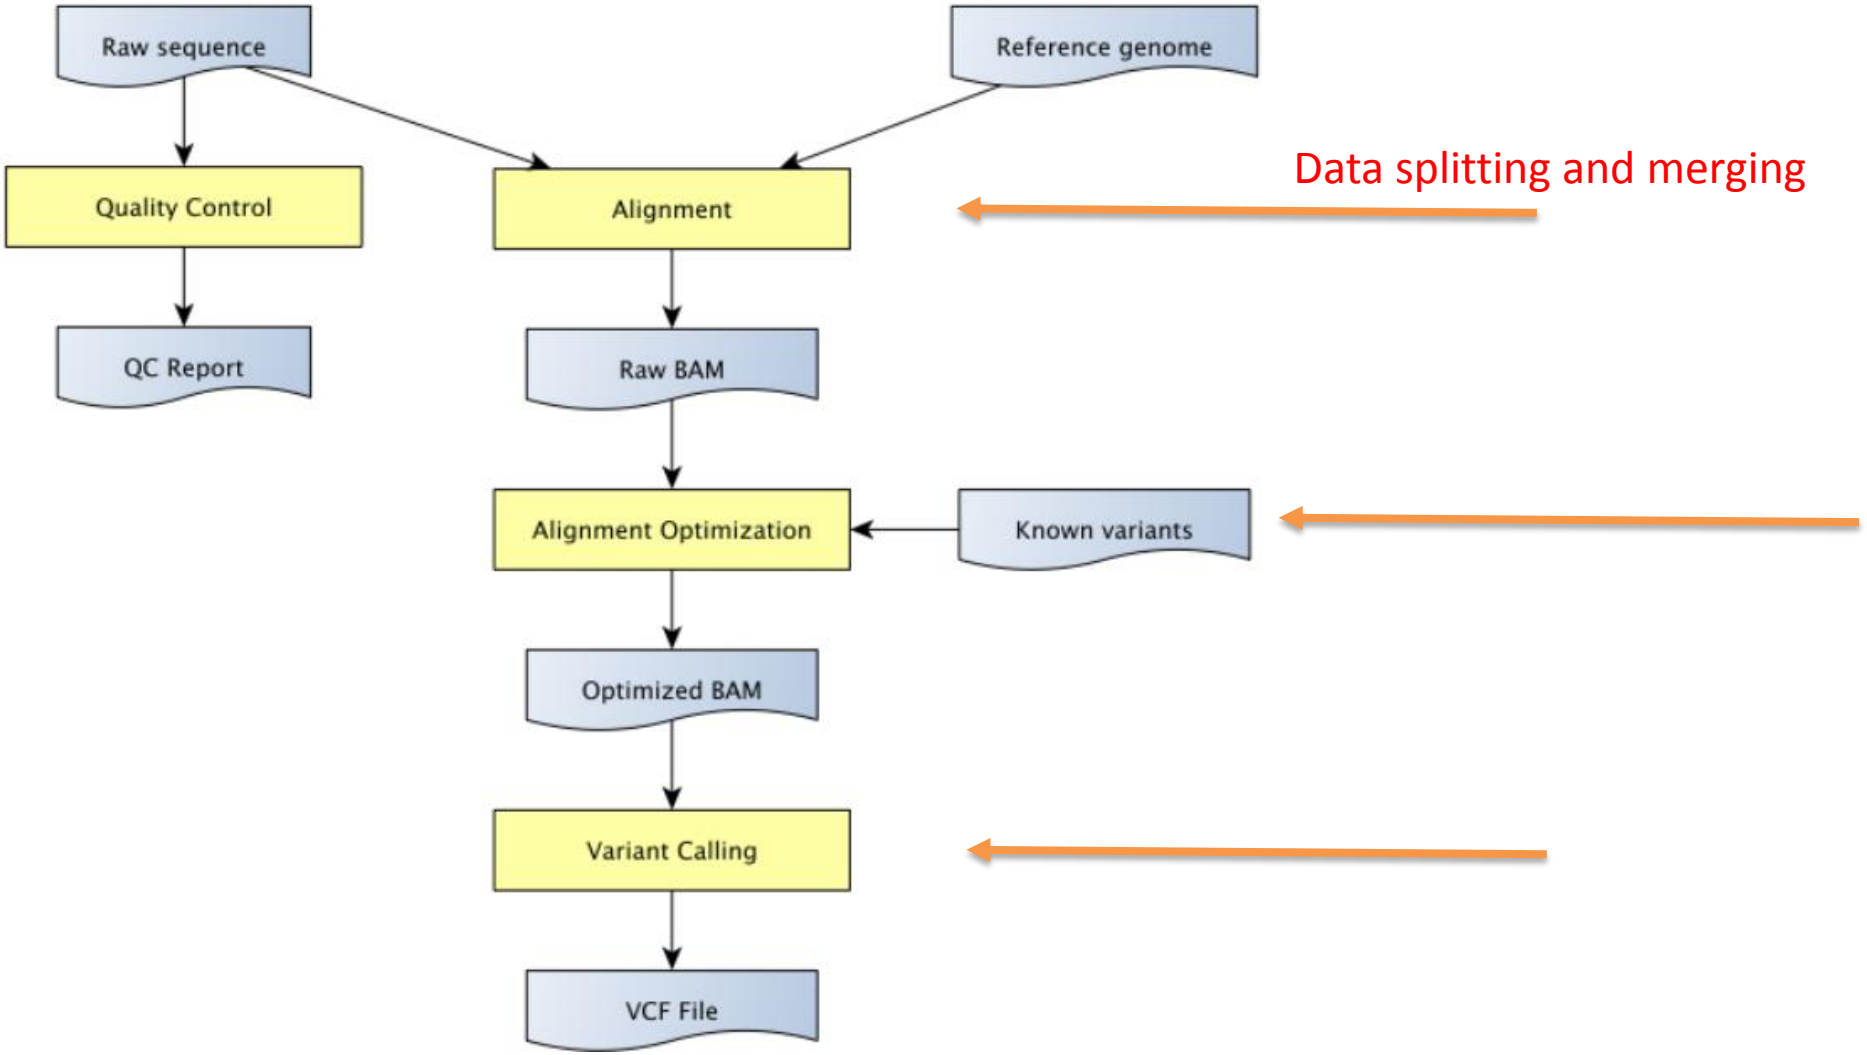

Data splitting and merging

How to handle data-parallel workflows?

Our “quick & dirty” solution: write a set of configurable scripts that handle job submissions, data dependencies and error trapping

What we really need is a workflow manager that understands data-parallel computing requirements

- Handles data chunking and merging
- Interfaces with job submission systems (PBS/Torque, LSF, SGE etc.)
- Understands data-level dependencies
- Performs error trapping and checkpointing and automates restarting of failed workflows

Another challenge

The IGB cluster is used by many people at the same time, and most are doing genomic analyses

While the GPFS system provides high nominal data throughput, it can slow to a crawl when heavily used

Slow I/O translates into severely degraded system performance

High-Throughput Computing (HTC) for Computational Genomics

"Traditional" HPC facilities can be leveraged to do computational genomics, but this takes some effort and technical savvy

There is a real need to design and build HTC facilities, designed specifically to support computational genomics

- These would look significantly different from existing HPC centers, and should be managed by a team including bioinformaticians
- In the US, no one seems to be willing to pay for this

A Notional HTC Architecture

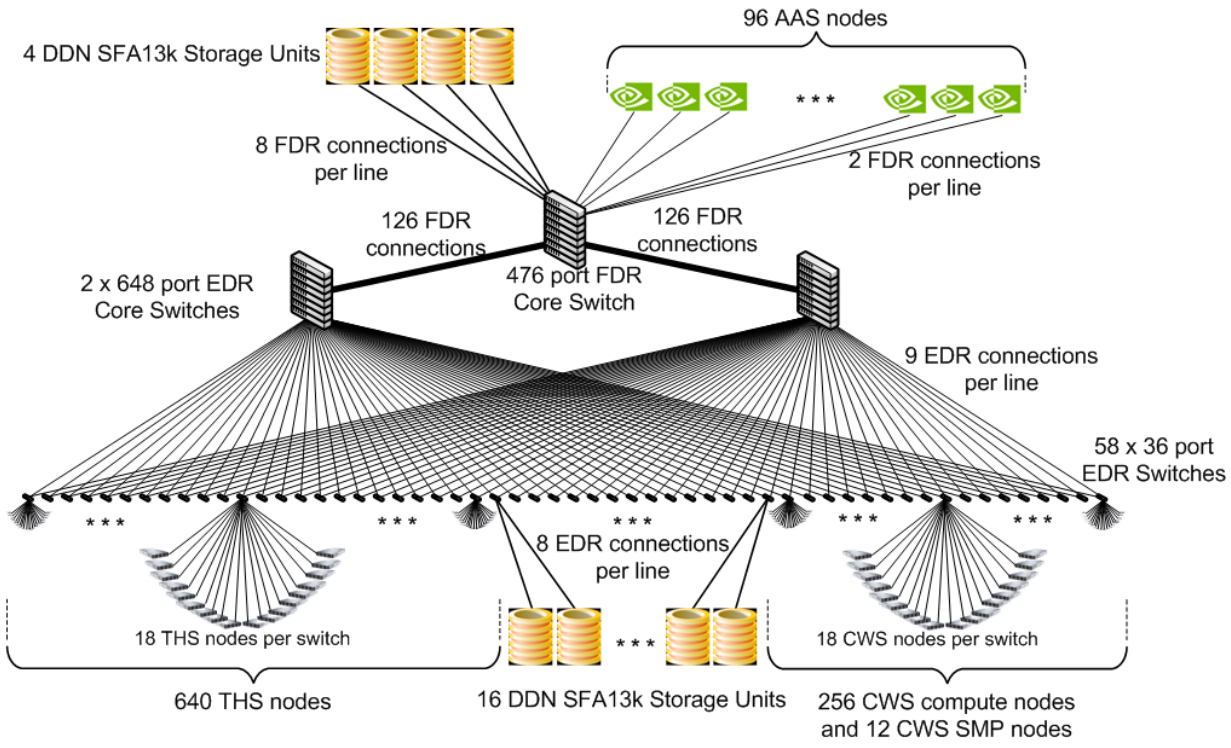

Four components:

- THS: traditional HPC cluster with high-speed interconnect (could be reduced / removed)
- AAS: accelerated applications subsystem with GPUs
- CWS: complex data-intensive workflows subsystem, including large memory (≥ 2 TB RAM) nodes
- HDS: high-performance data subsystem with FPGA-accelerated database server

Optimized for:

- Fast I/O for both large files and many small files
- Workflows that pass data from one application to another
- Embarrassingly parallel execution
- Built-in Map/Reduce capability

Innovations Required for a True HTC System

Application level

- » Software that leverages the hardware architecture efficiently (node and system level parallelism, GPU/FPGA components)

Data level

- » Fast highly parallel storage, resilient to concurrent access and large number of small files
- » New data structures to represent the complexity of genomic data

System level

- » Workflow manager designed for complex data-parallel pipelines

The HPCBio team

Carver Biotechnology Center

- Chris Fields
- Radhika Khetani
- George Gong
- Ravi Kiran Donthu
- Jenny Drnevich Zadeh
- Janet Chang

IGB CNRG

- Dan Davidson
- David Slater
- Nevo Band

NCSA / IGB

- Gloria Rendon
- Luda Mainzer

Free electron

- Reza Farivar (IGB and OTM)