

Yu-Ting Chen¹, Jason Cong¹, Jie Lei¹, Sen Li¹, Myron Peto², Paul Spellman², Peng Wei¹, and Peipei Zhou¹

Contact: ytchen@cs.ucla.edu

¹ Computer Science Department, UCLA, USA

² Dept. of Molecular and Medical Genetics, Oregon Health State University, USA

Available at: <https://github.com/ytchen0323/cloud-scale-bwamem>

Motivation: Build a Fast and Scalable Read Aligner for WGS Data

- Target: whole-genome sequencing
 - A huge amount of data (500M ~ 1B reads for 30x coverage)
 - Pair-end sequencing
 - Goal: improve the speed of whole-genome sequencing
- Limitation of the state-of-the-art aligners
 - BWA-MEM takes about 10 hours to align WGS data
 - Multi-threaded parallelization in a single server
 - The speed is limited by the single-node computation power and I/O bandwidth when data size is huge
- Proposed tool: Cloud-Scale BWAMEM
 - Leverage the BWA-MEM algorithm
 - Exploit the enormous parallelism by using cloud infrastructures

Tool Highlight

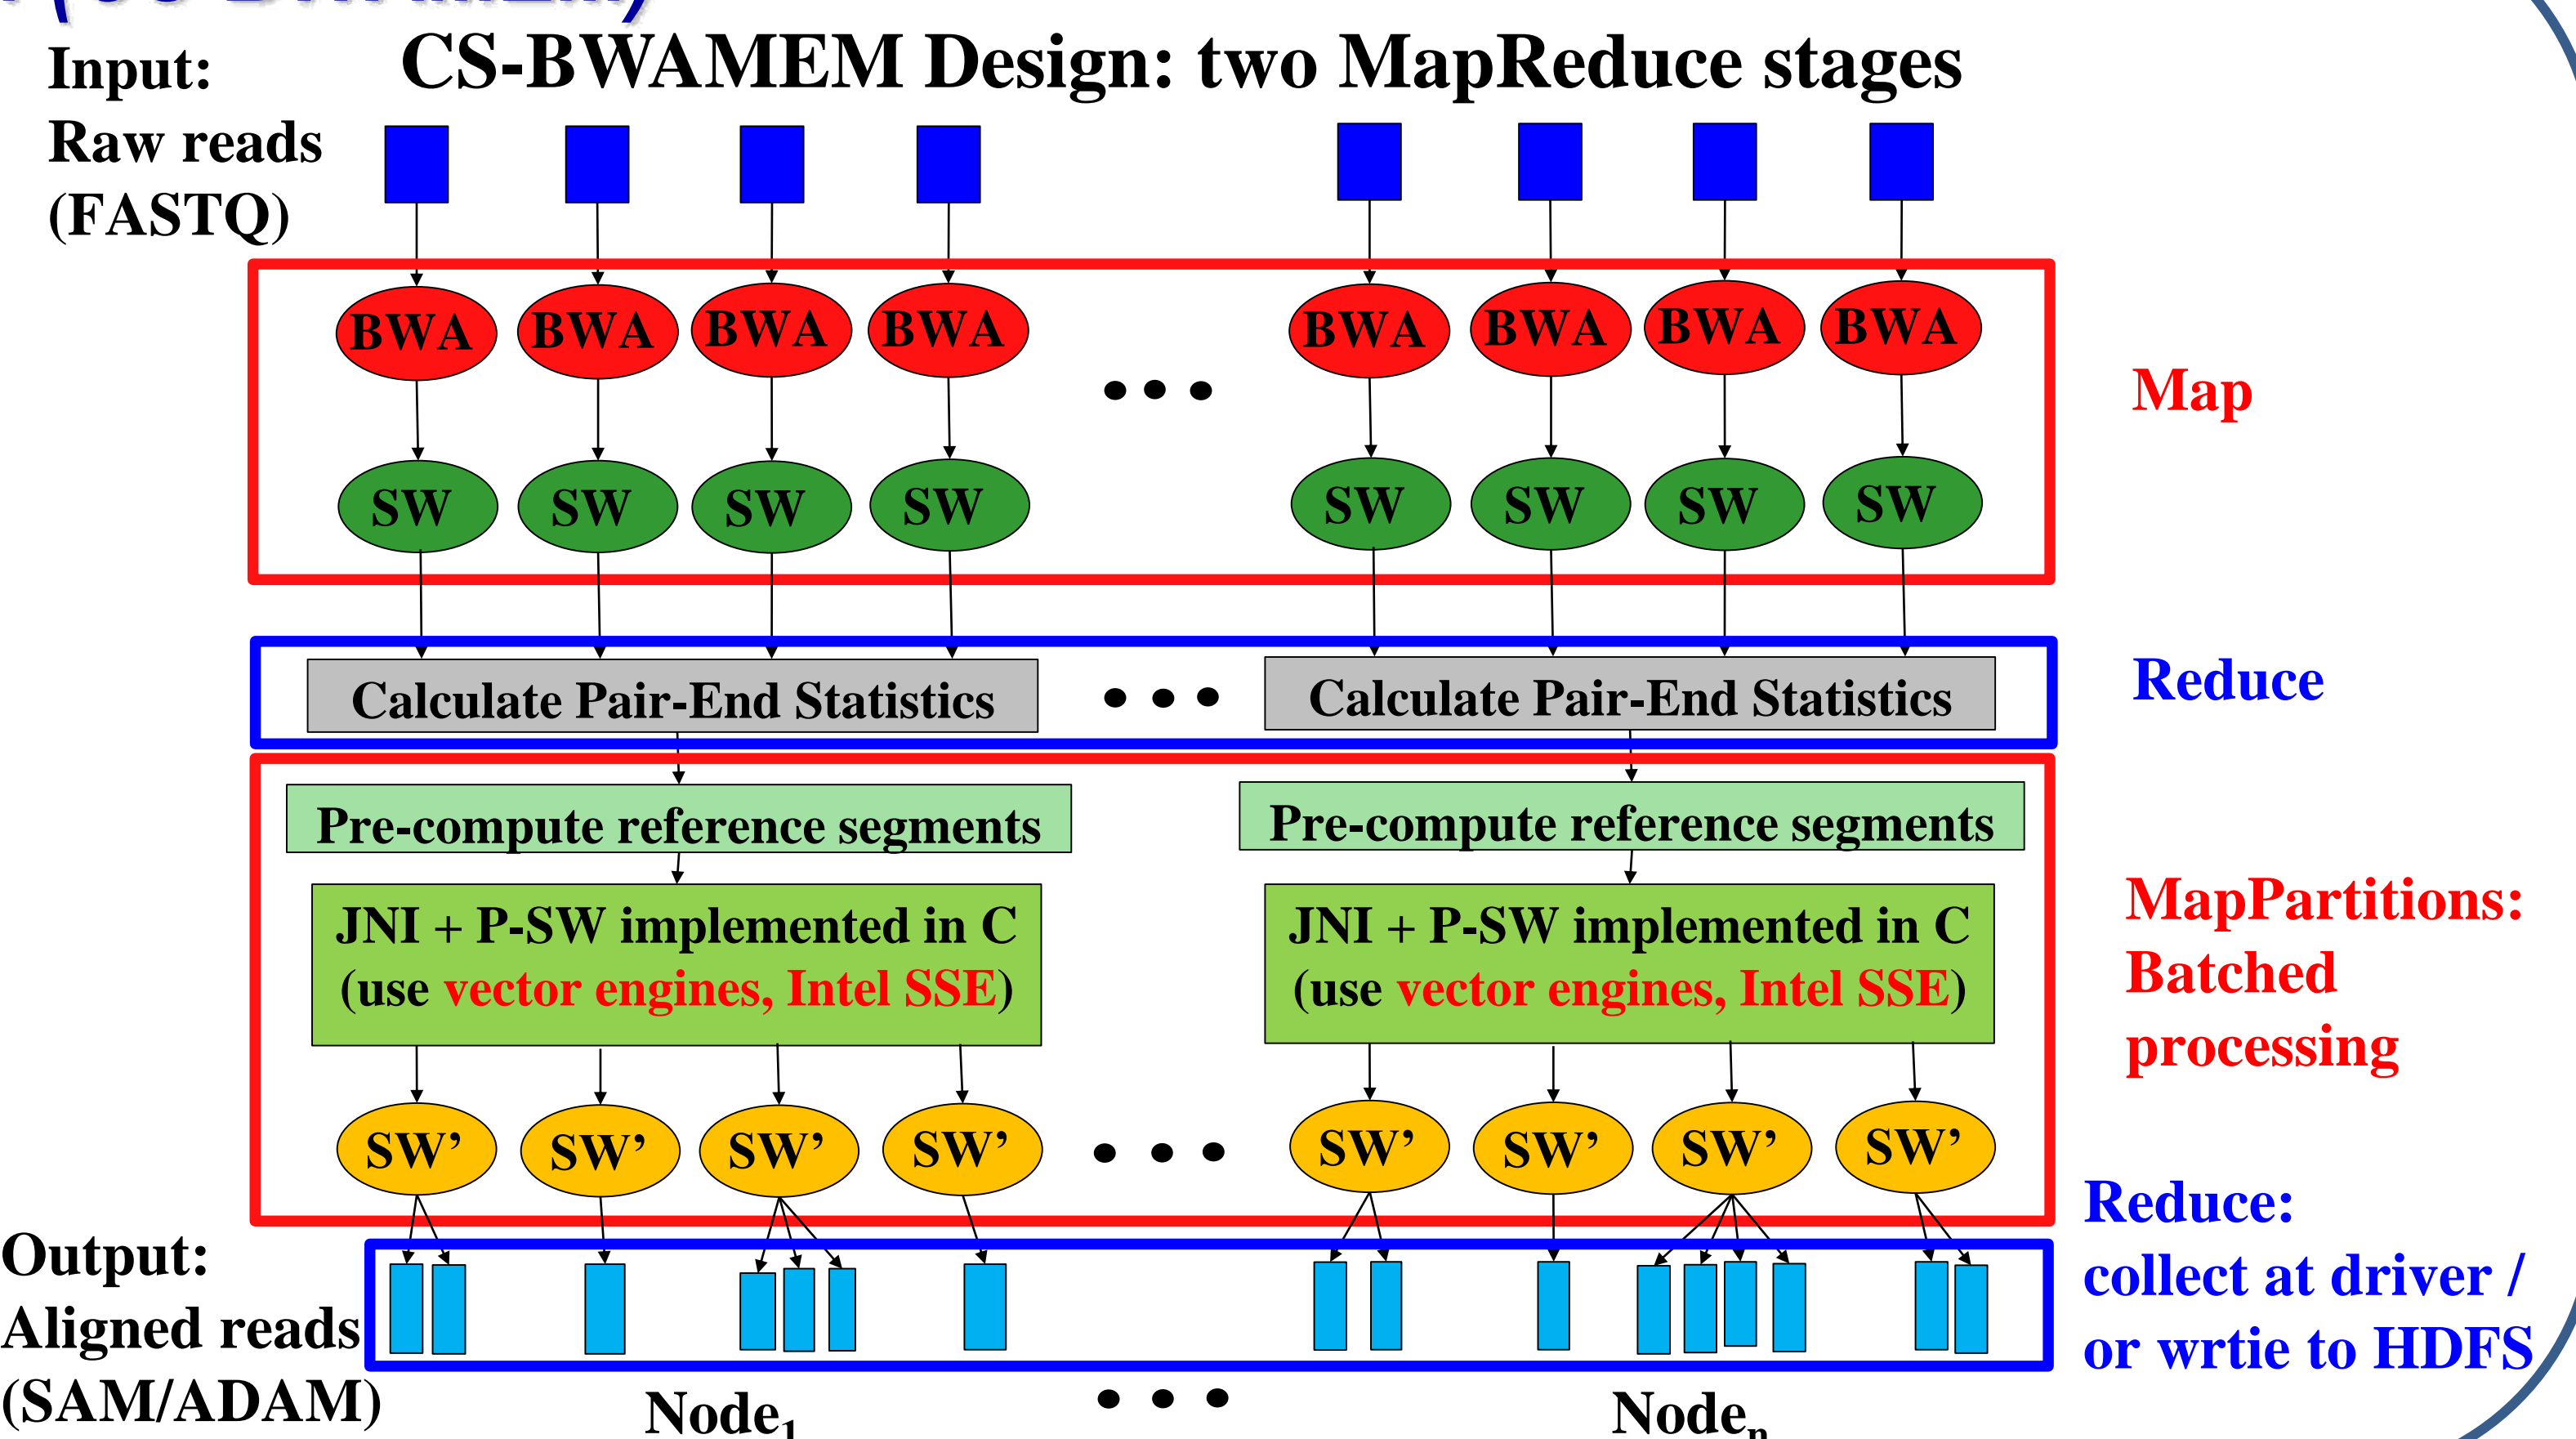
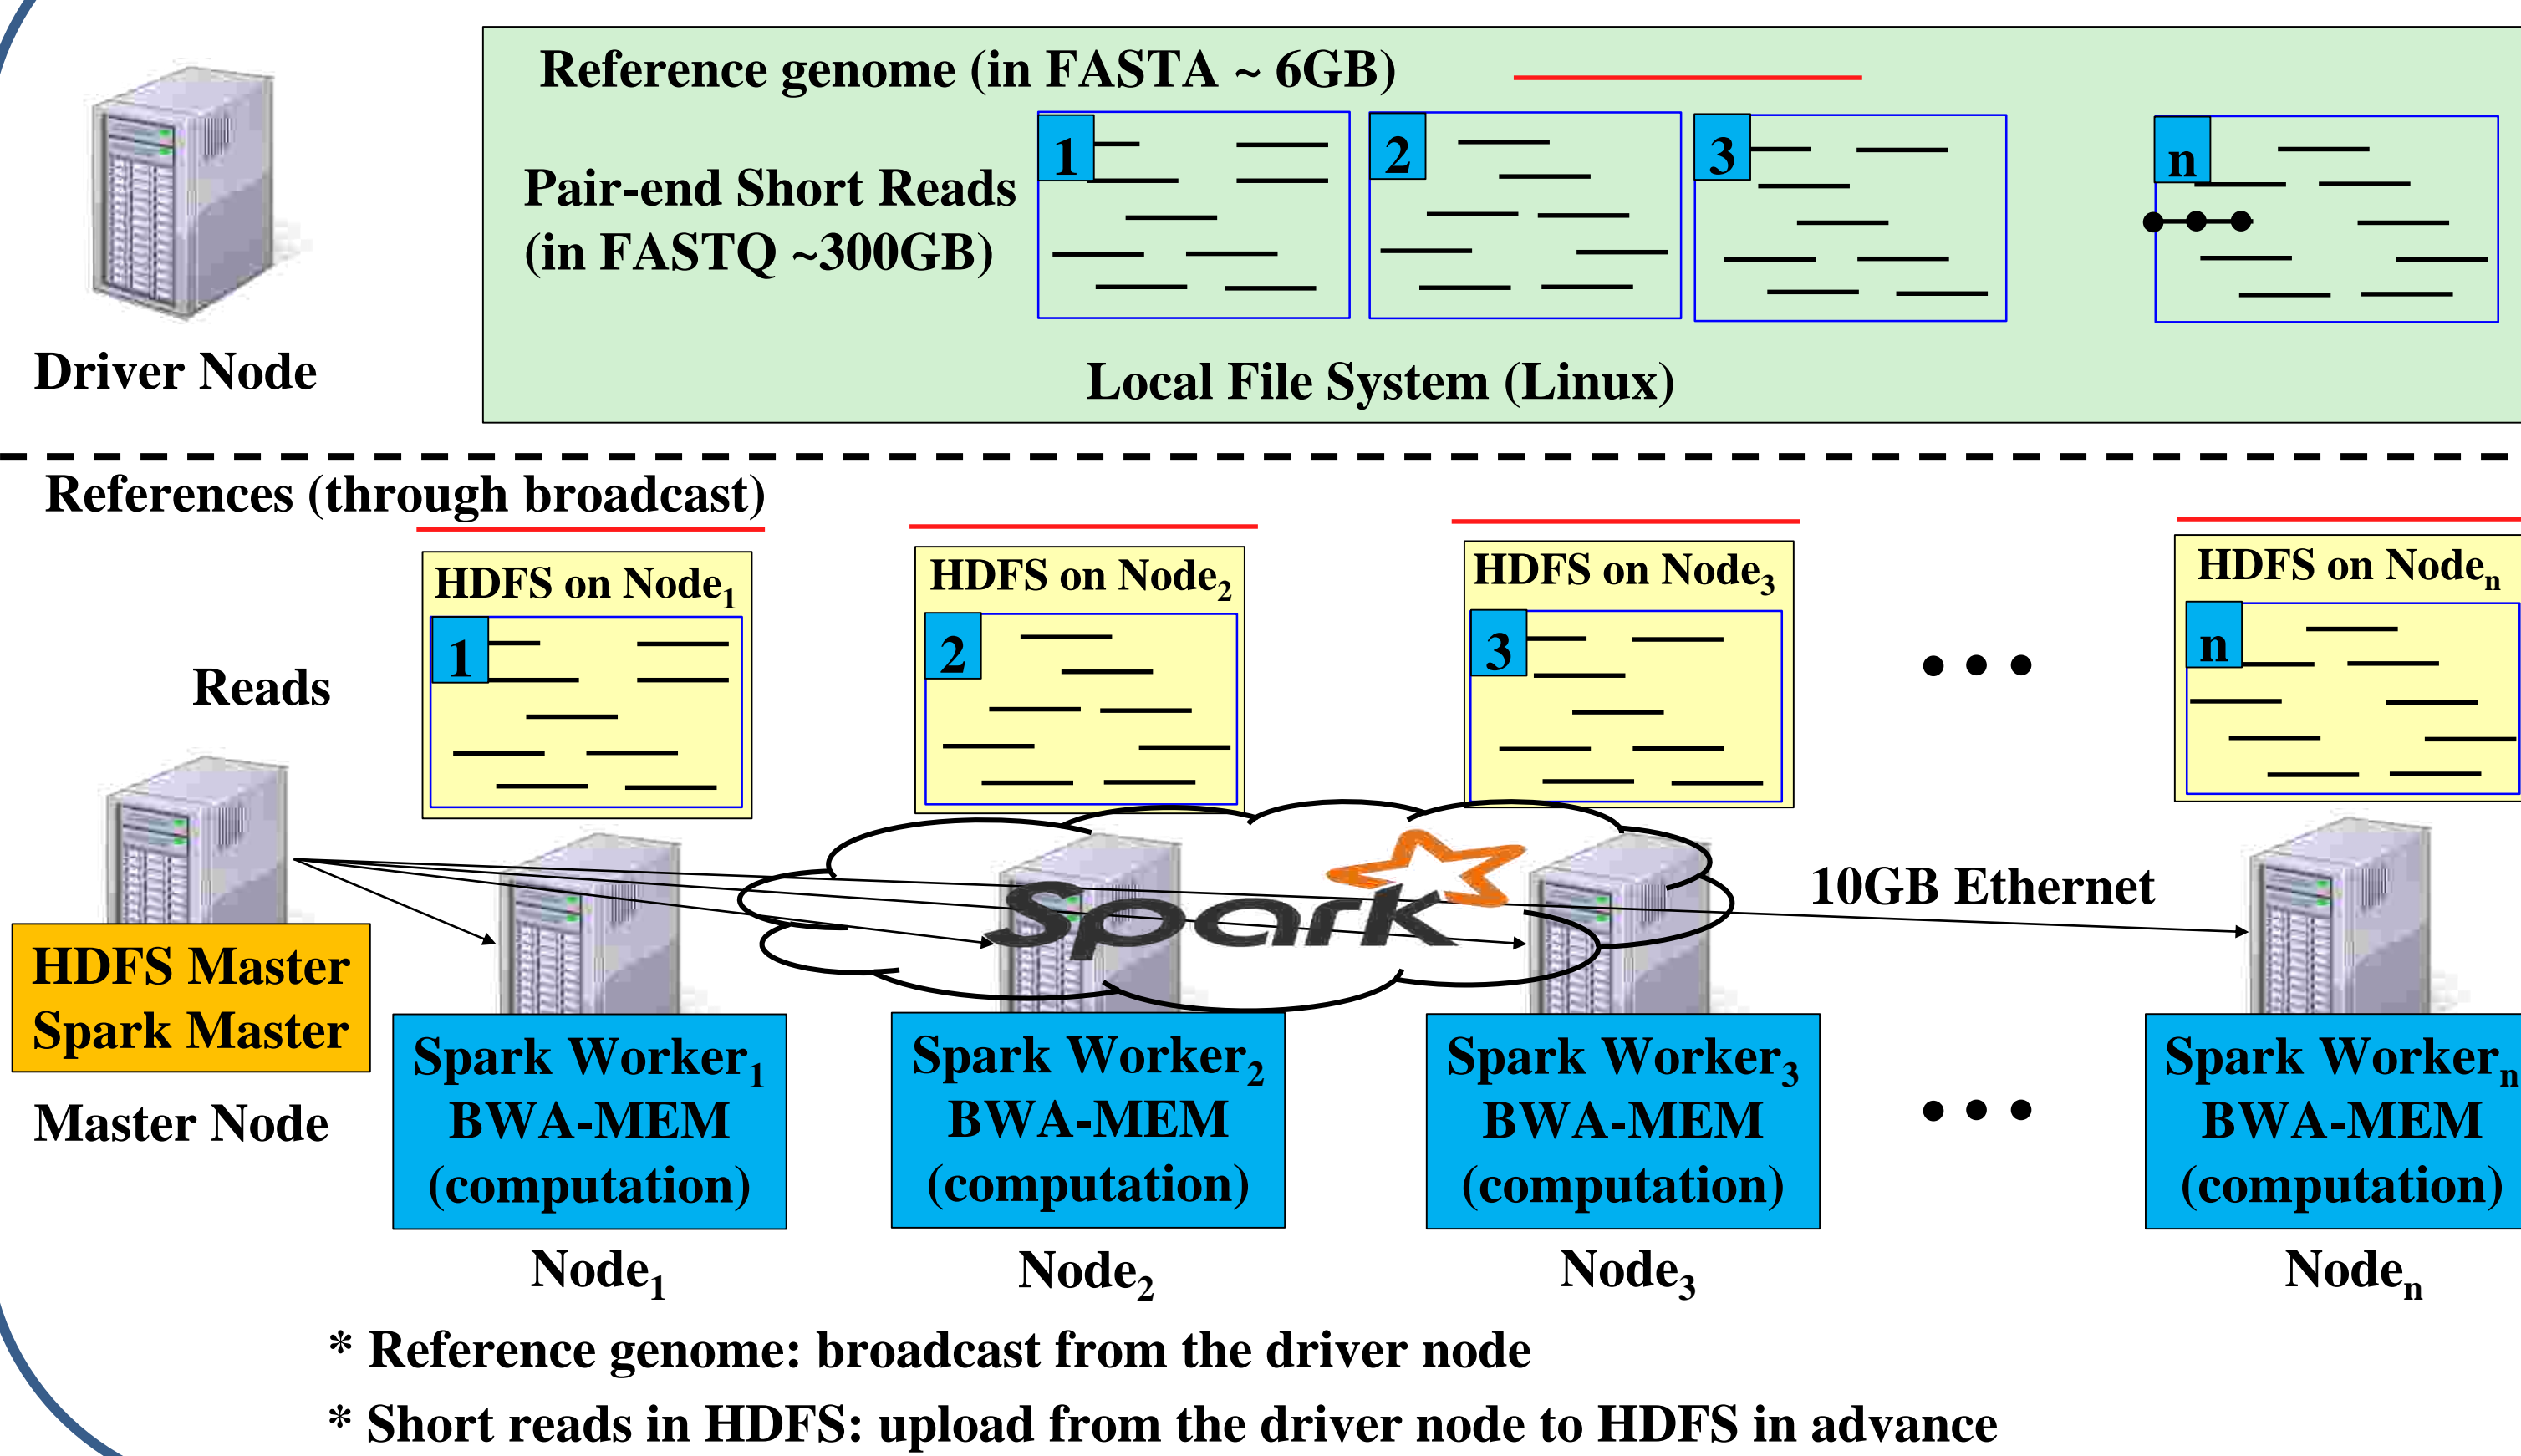
- Provide **scalable** speedup for read alignment
 - Users can choose an adequate number of nodes in their cluster based on their alignment performance target
- Speed of aligning a whole-genome sample
 - < 80 minutes: whole-genome data (30x coverage; ~300GB)
 - 25-node cluster with 300 cores
 - BWA-MEM: 9.5 hours
- Features
 - Support both pair-end and single-end alignment
 - Achieve similar quality to BWA-MEM
 - Input: FASTQ files
 - Output: SAM (single-node) or ADAM (cluster) format

Methods

- Use **MapReduce** programming model in a **cluster**
 - Most commonly used for big data analytics
 - Good for large-scale deployment – handle enormous parallelism of input reads
- Computation infrastructure: **Spark**
 - In-memory MapReduce system
 - Cache intermediate data in memory for future steps
 - Avoid unnecessary slow disk I/O accesses
- Storage infrastructure: **Hadoop distributed file system (HDFS)**
 - HDFS brings scalable I/O bandwidth since the disk I/O are linearly proportional to the size of a cluster
 - Store the FASTQ input in a distributed fashion
 - Spark can get data from HDFS before processing




Cloud-Scale BWAMEM (CS-BWAMEM)



Performance

- Input: whole-genome data sample (30x coverage; about 300GB FASTA files)
- 7x speedup over BWA-MEM
 - Can align whole-genome data **within 80 minutes**
 - User can adjust the cluster size based on their performance demand


BWA-MEM vs. CS-BWAMEM



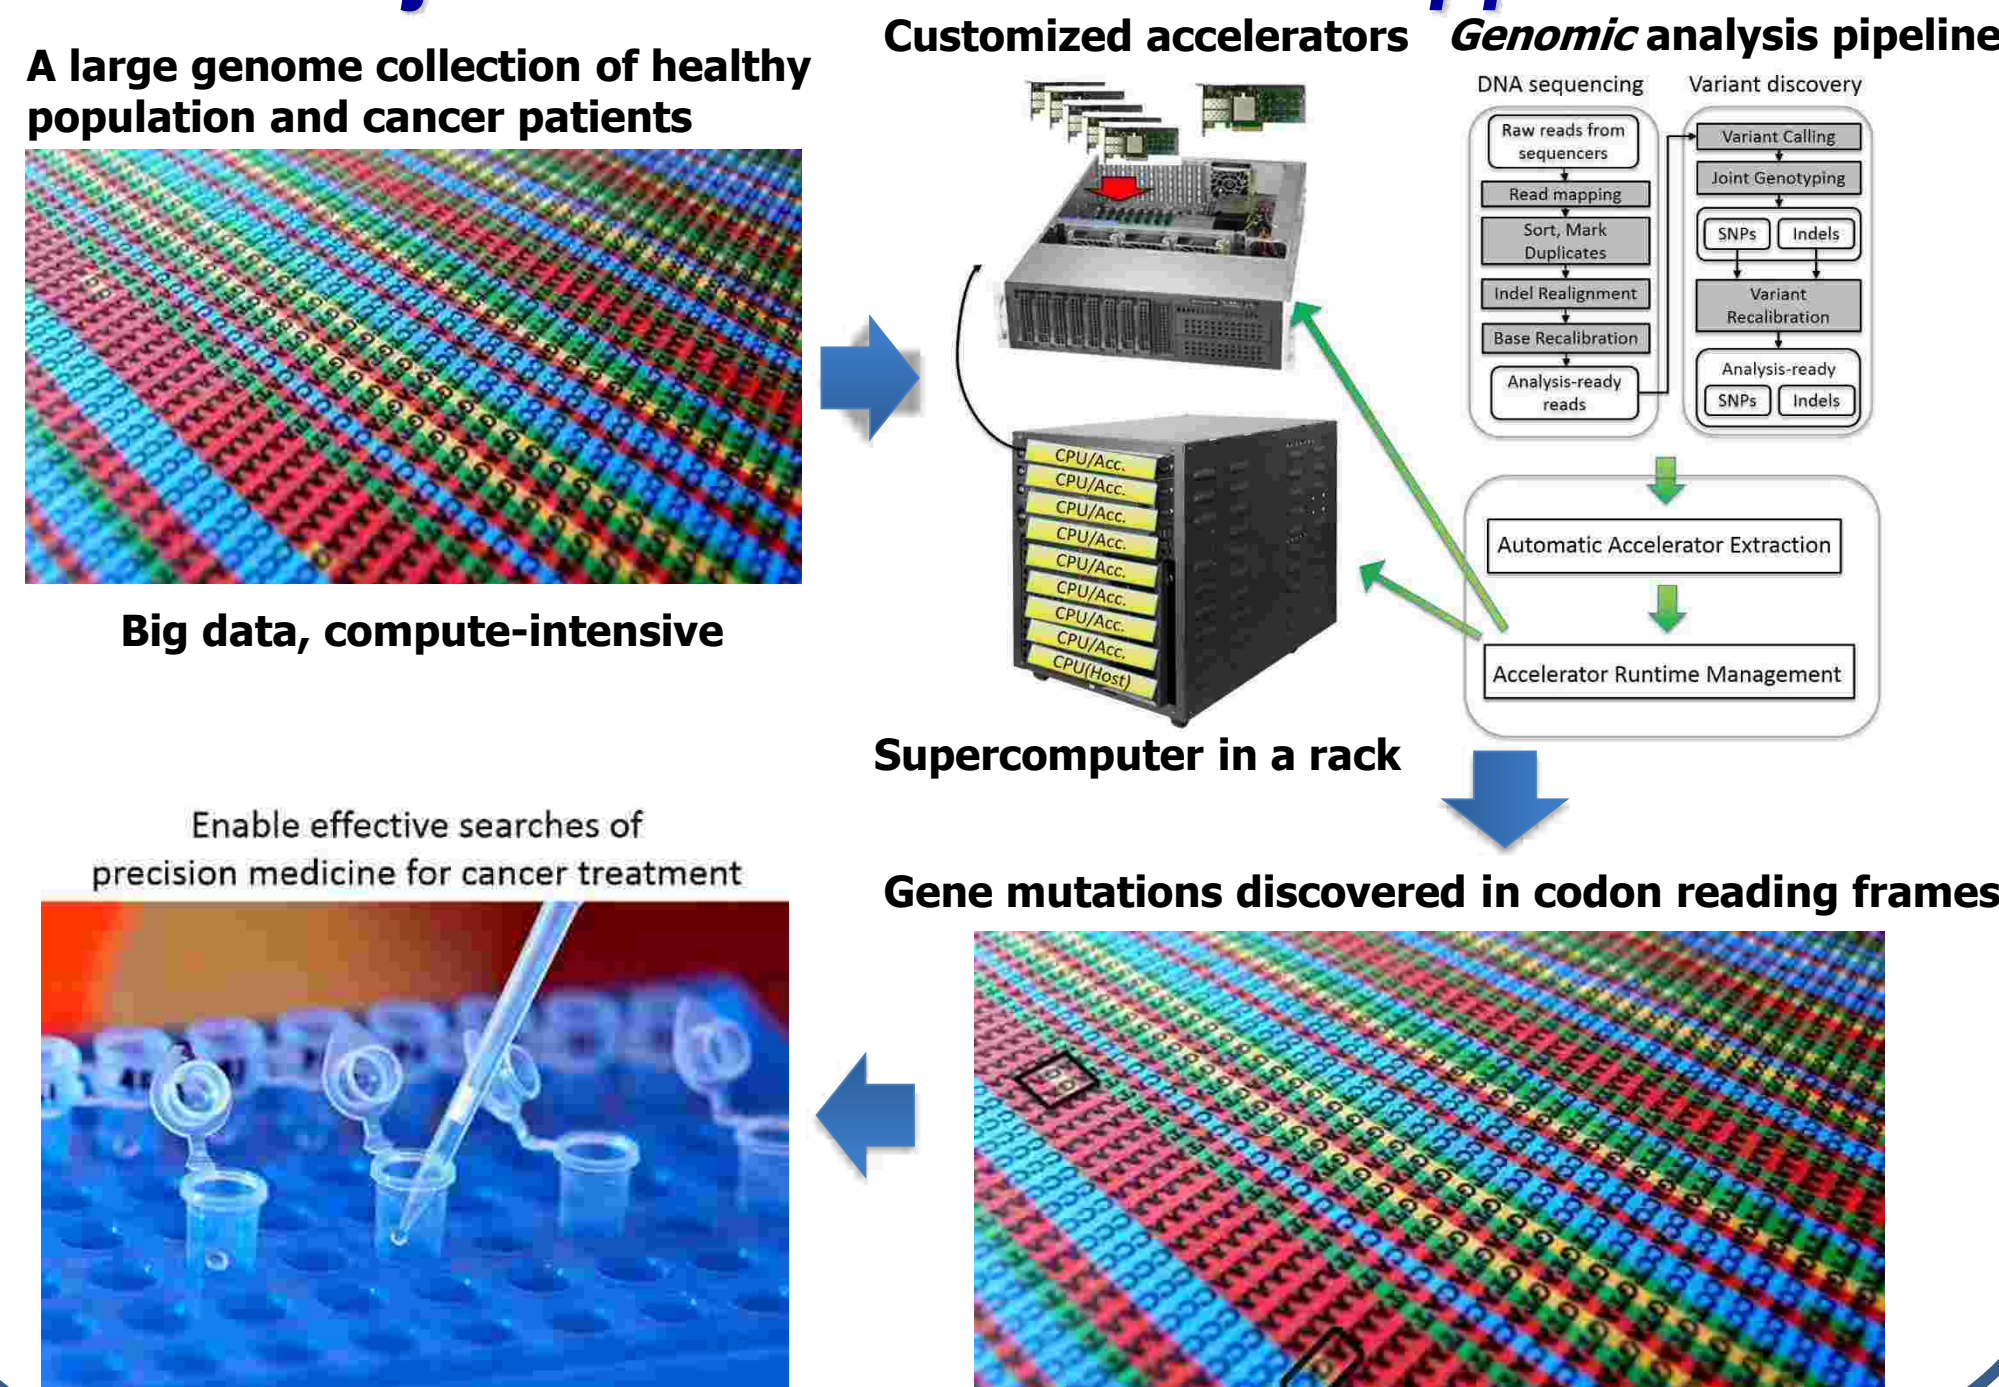
Comparison with BWA-MEM

- Flow: CS-BWAMEM / BWA-MEM -> sort -> markduplicate -> indelrealignment -> baserecalibration
- Detailed analysis on two cancer exome samples
 - Buffy sample and primary tumor sample
- Almost same mapping quality to BWA-MEM
- The numbers of mapped reads:
 - CS-BWAMEM vs. BWA-MEM
 - Buffy sample: 99.59% vs. 99.59%
 - Primary sample: 99.28% vs. 99.28%
- Difference on total reads
 - CS-BWAMEM vs. BWA-MEM
 - Buffy sample: 0.006%
 - Primary sample: 0.04%

Alignment Quality




Our Project & Cancer Genome Applications



Cluster Deployment

One master / driver node
10GbE switch
25 worker nodes



- Server node setting
 - Intel Xeon server
 - Two E5-2620 v3 CPUs
 - 64GB DDR3/4 RAM
 - 10GBE NIC
- Software infrastructure
 - Spark 1.3.1 (v0.9 – v1.3.0 tested)
 - Hadoop 2.5.2 (v2.4.1 tested)
- Hardware acceleration
 - On-going project
 - PCIe-based FPGA board
 - Ex: Smith-Waterman / Burrows-Wheeler transform kernels

Collaborations

- Test our aligner on the data from our collaborators
- Oregon Health State University (OHSU)
 - Spellman Lab
 - Application: cancer genomics and precision medicine
- UCLA
 - Coppola Lab
 - Application: neuroscience
 - Neurodegenerative conditions, including Alzheimer's Disease (AD), Frontotemporal Dementia (FTD), and Progressive Supranuclear Palsy
 - Fan Lab
- University of Michigan, Ann Arbor
 - University of Michigan comprehensive cancer center (UMCCC)
 - Application:
 - Motility based cell selection for understanding cancer metastasis.