

Analysis of high-throughput sequencing data with HPC

Frédéric Jarlier



- **Plan**

- Presentation
- Context
- The alignment
- The sorting
- Conclusion
- Questions

- **Presentation**

Institut Curie is the cancer foundation of public interest founded in 1909

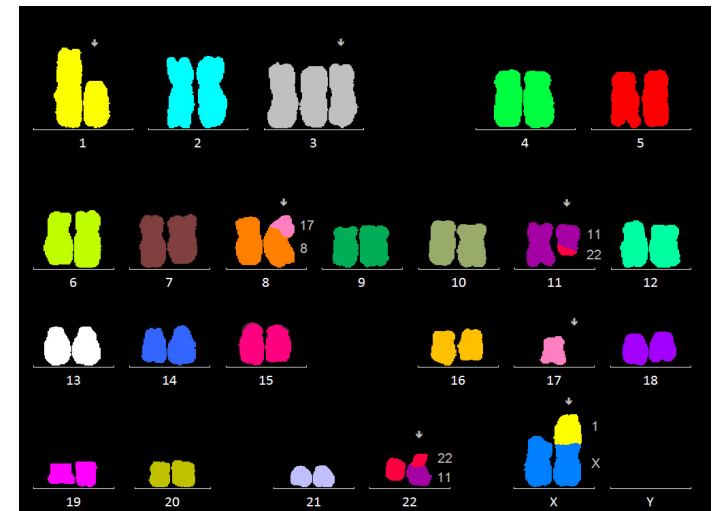
Specialized in research in oncology and patient care

Institut Curie has 3,400 researchers, physicians, clinicians, technicians and administrative staff




- **Context**

- Understanding the cancer is understanding the genome
- The cancer is a gene disease
- To decipher the genome we use sequencing





CK Rocha et al., Molecular Cytogenetics 2011


• Data volume constraints



Decreasing cost of sequencing technology



Precision medicine

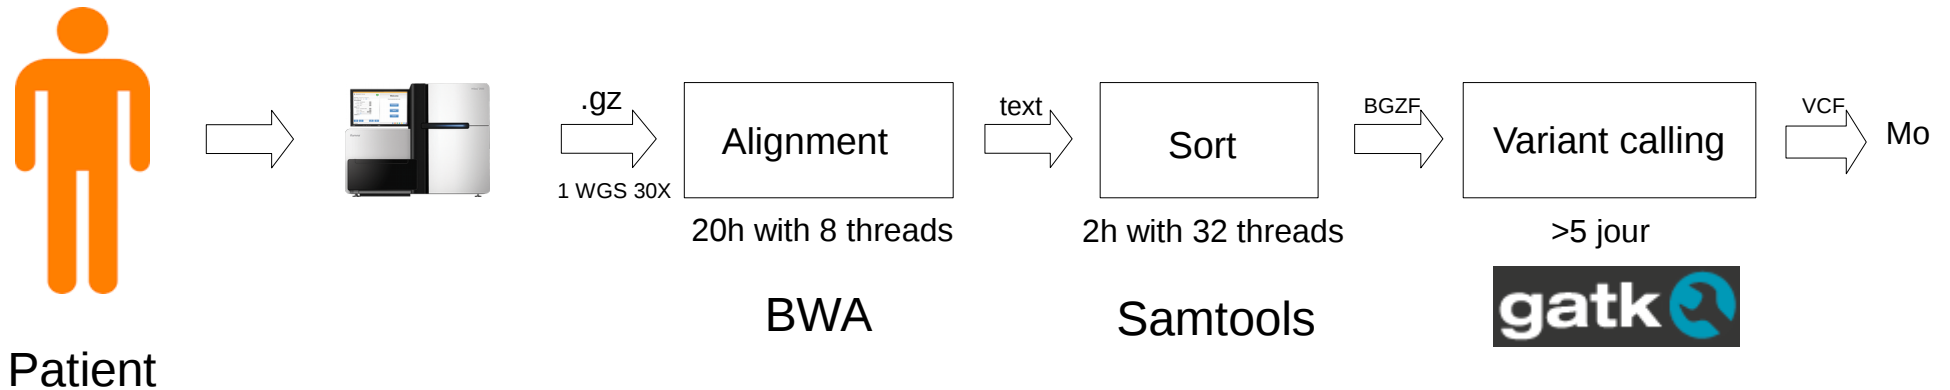


Increasing volume

Zachari S. et al PlosBio (2015)

- **Actual pipeline**

- Actual whole genome sequencing pipeline



Constit (30X) + Tumor (60X) = 1TB / patient

Analysis > 1 week

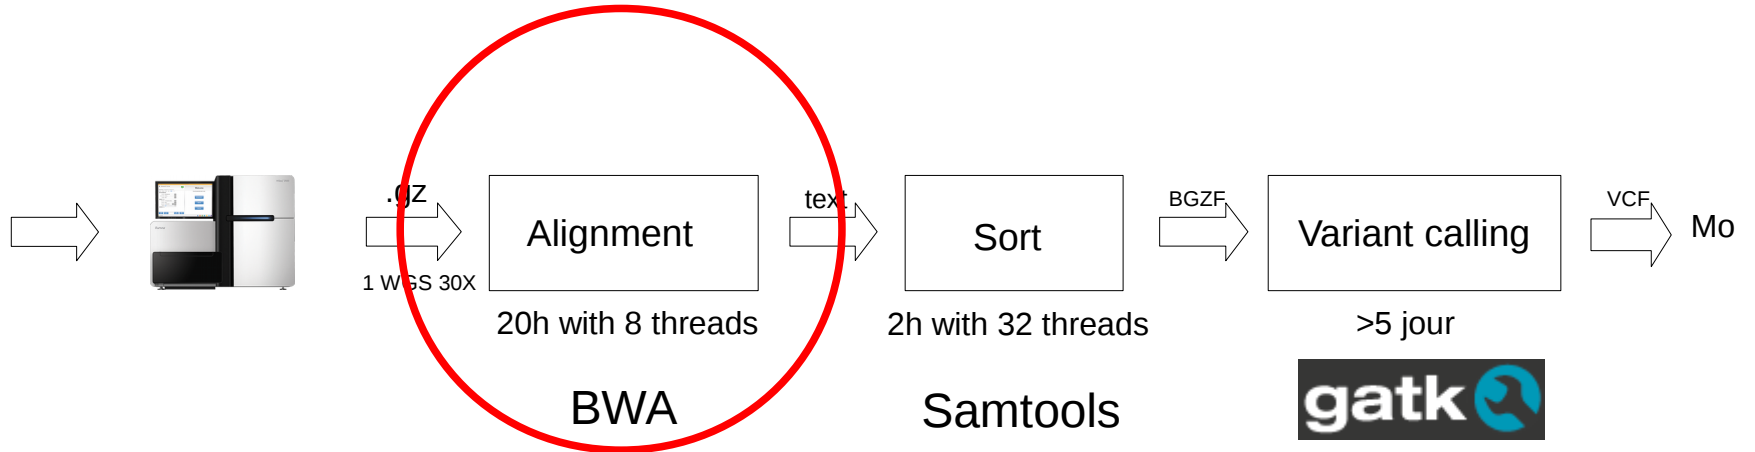
how do we make the pipeline scale?

Alignment

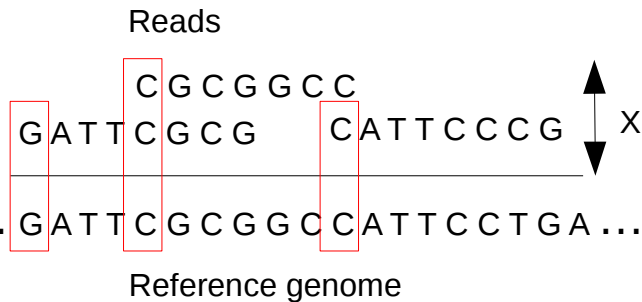
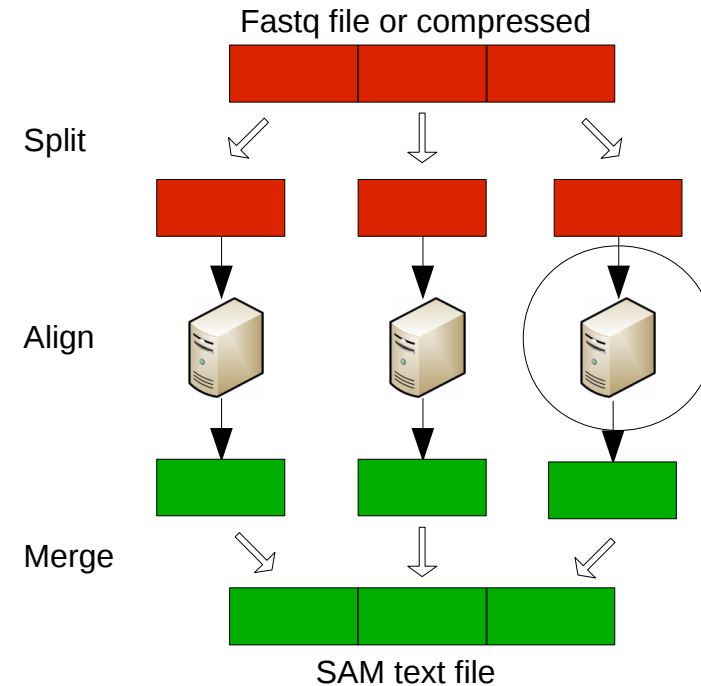
- Search coordinates of reads on reference genome

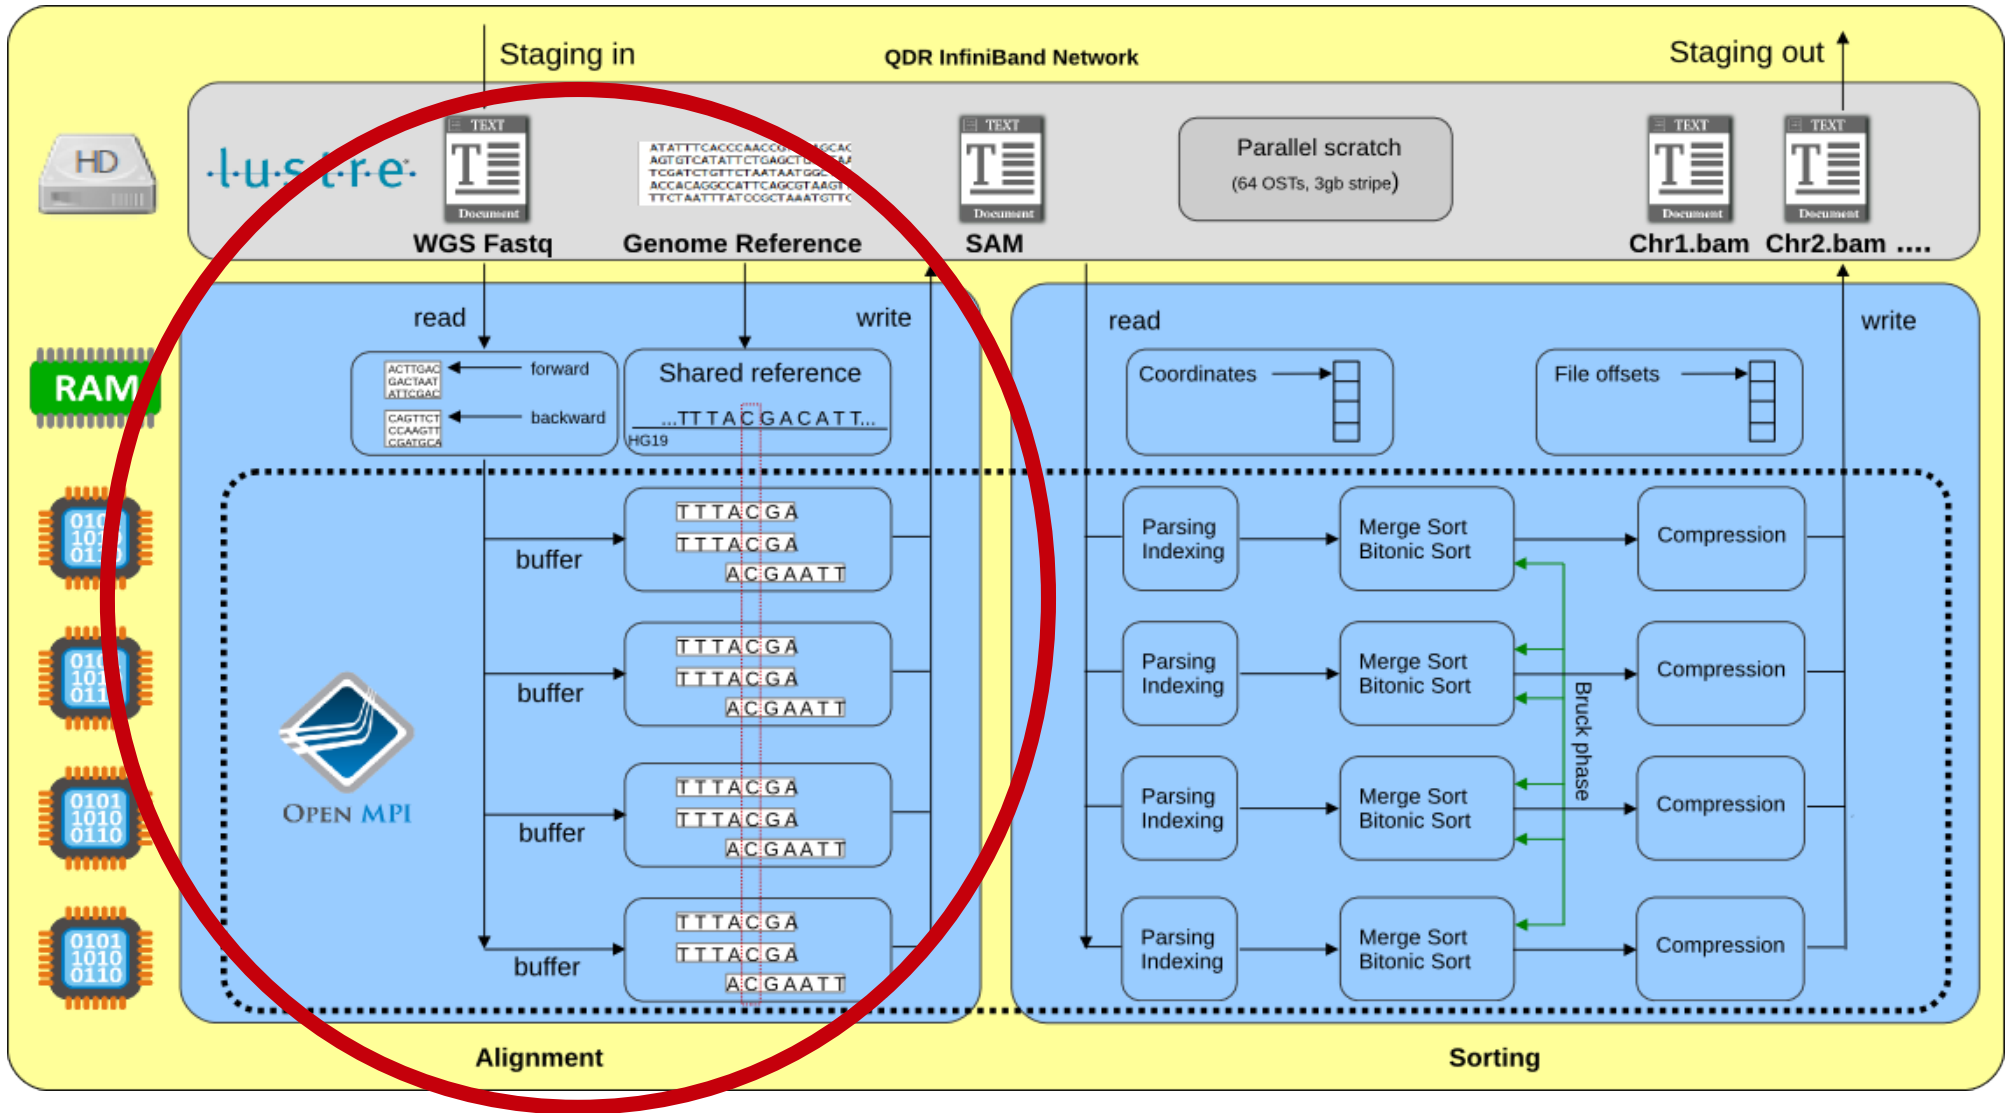


Patient

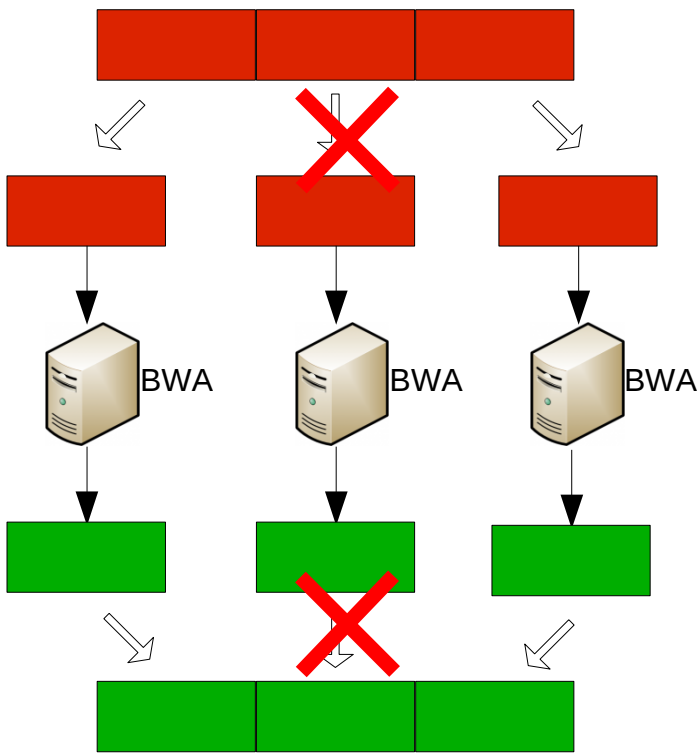


Embarassingly parallel



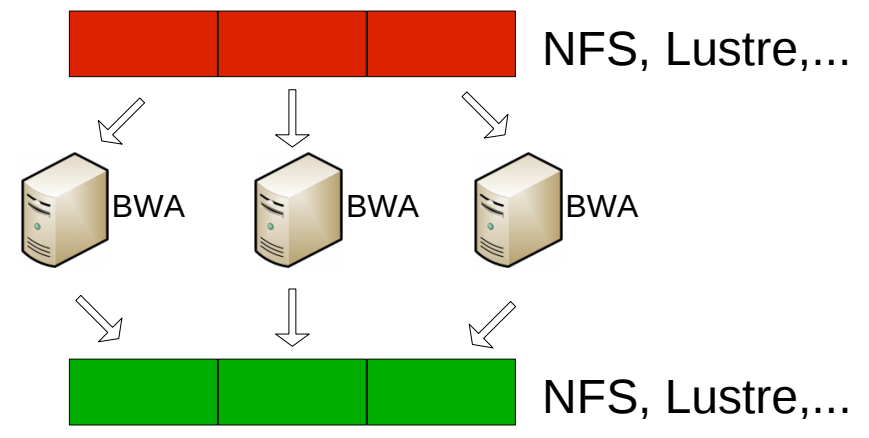


<https://github.com/InstitutCurie/QUASART>

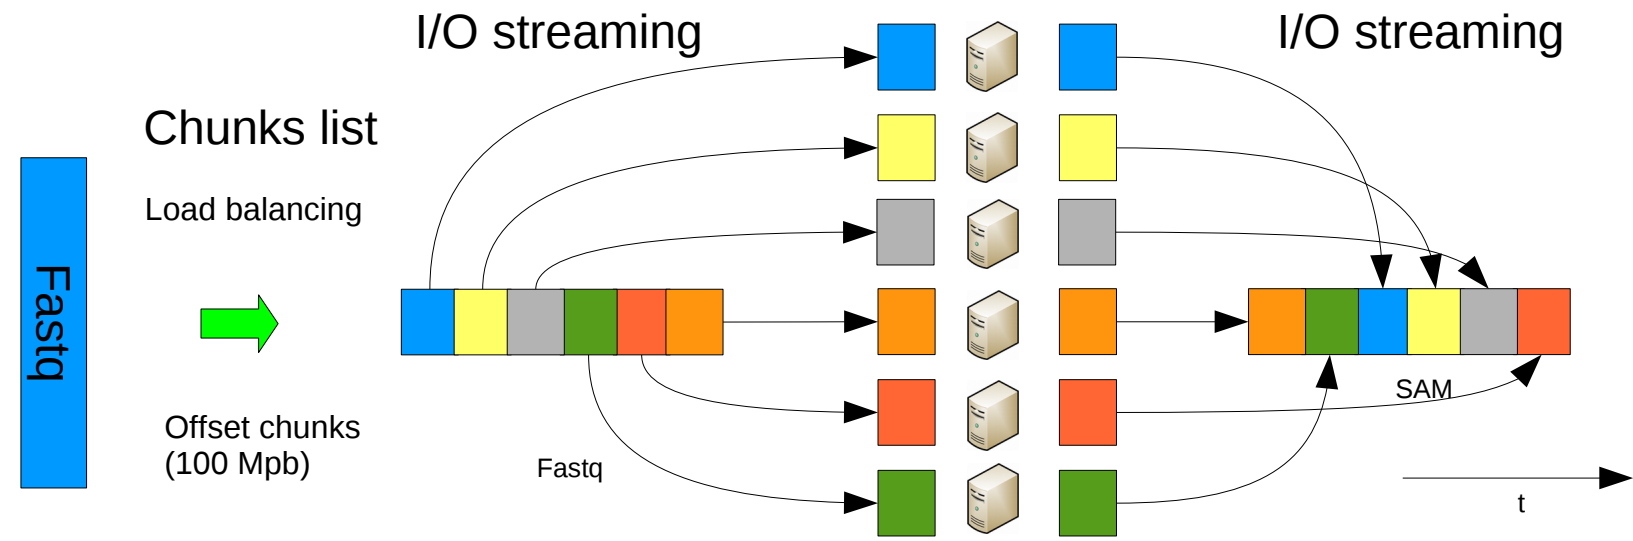


Eliminate split and cat

Shared File pointer operations (MPI)



Alignment



- **Tested Architectures**

180 dual processor nodes (Intel Sandy Bridge E5-2680, 2.7 GHz, 8 cores)

128 GB of memory per node, i.e. a total of 2,880 cores (Bull),

QDR Infiny Band (36 GB/s)

Lustre filesystem




Intel Swindon Wildcat

14 dual nodes (Xeon® E5-2697 v3 - 2.6 GHz 14 Core)




128 GB DDR4 per node

Intel® True Scale™ 12300

Lustre filesystem version: Intel Enterprise Edition for Lustre (175T)



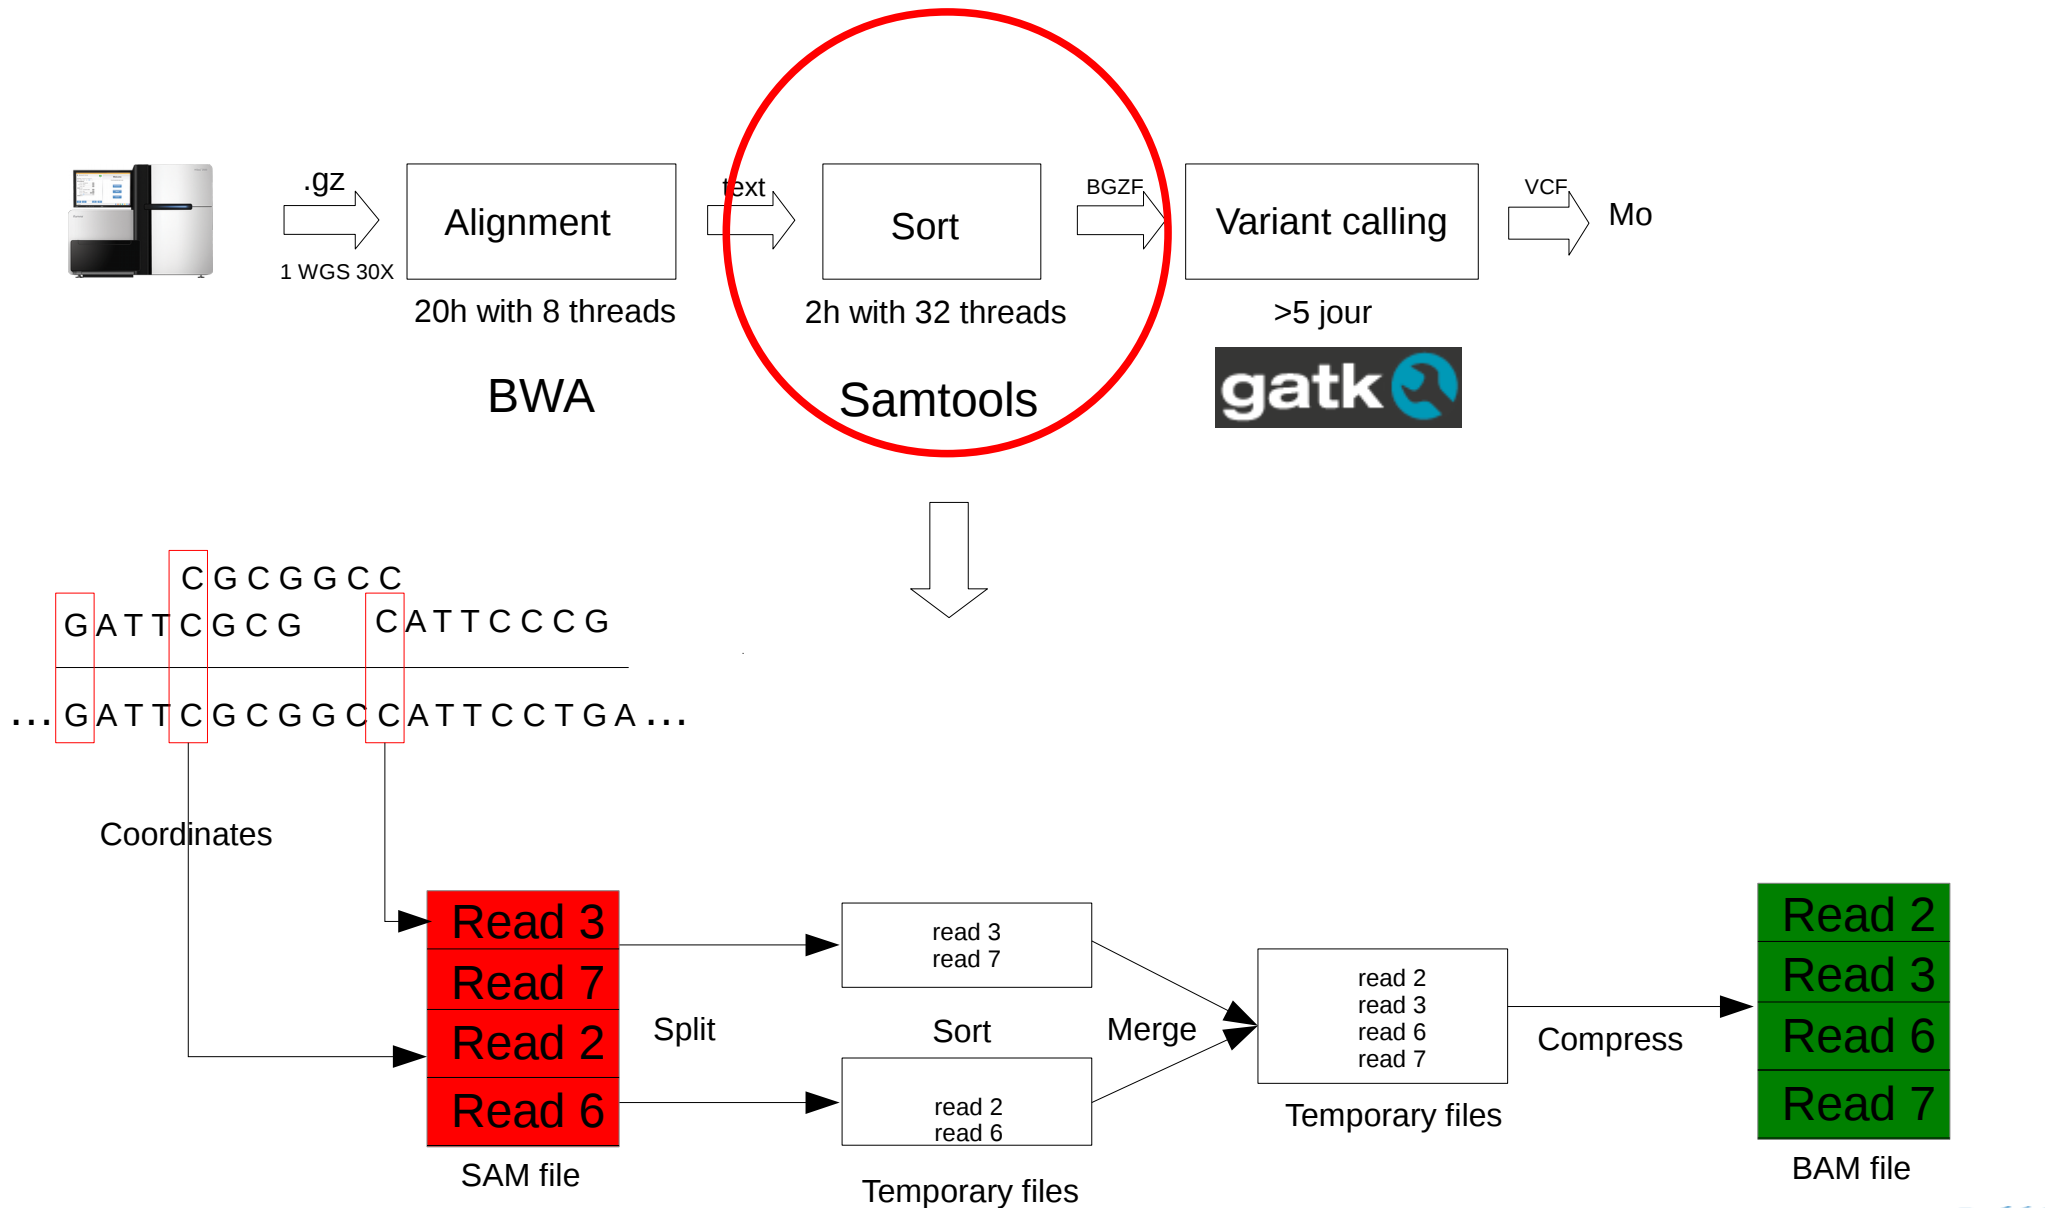
Wall Times and speed-up of the parallel alignment 93X WG pe

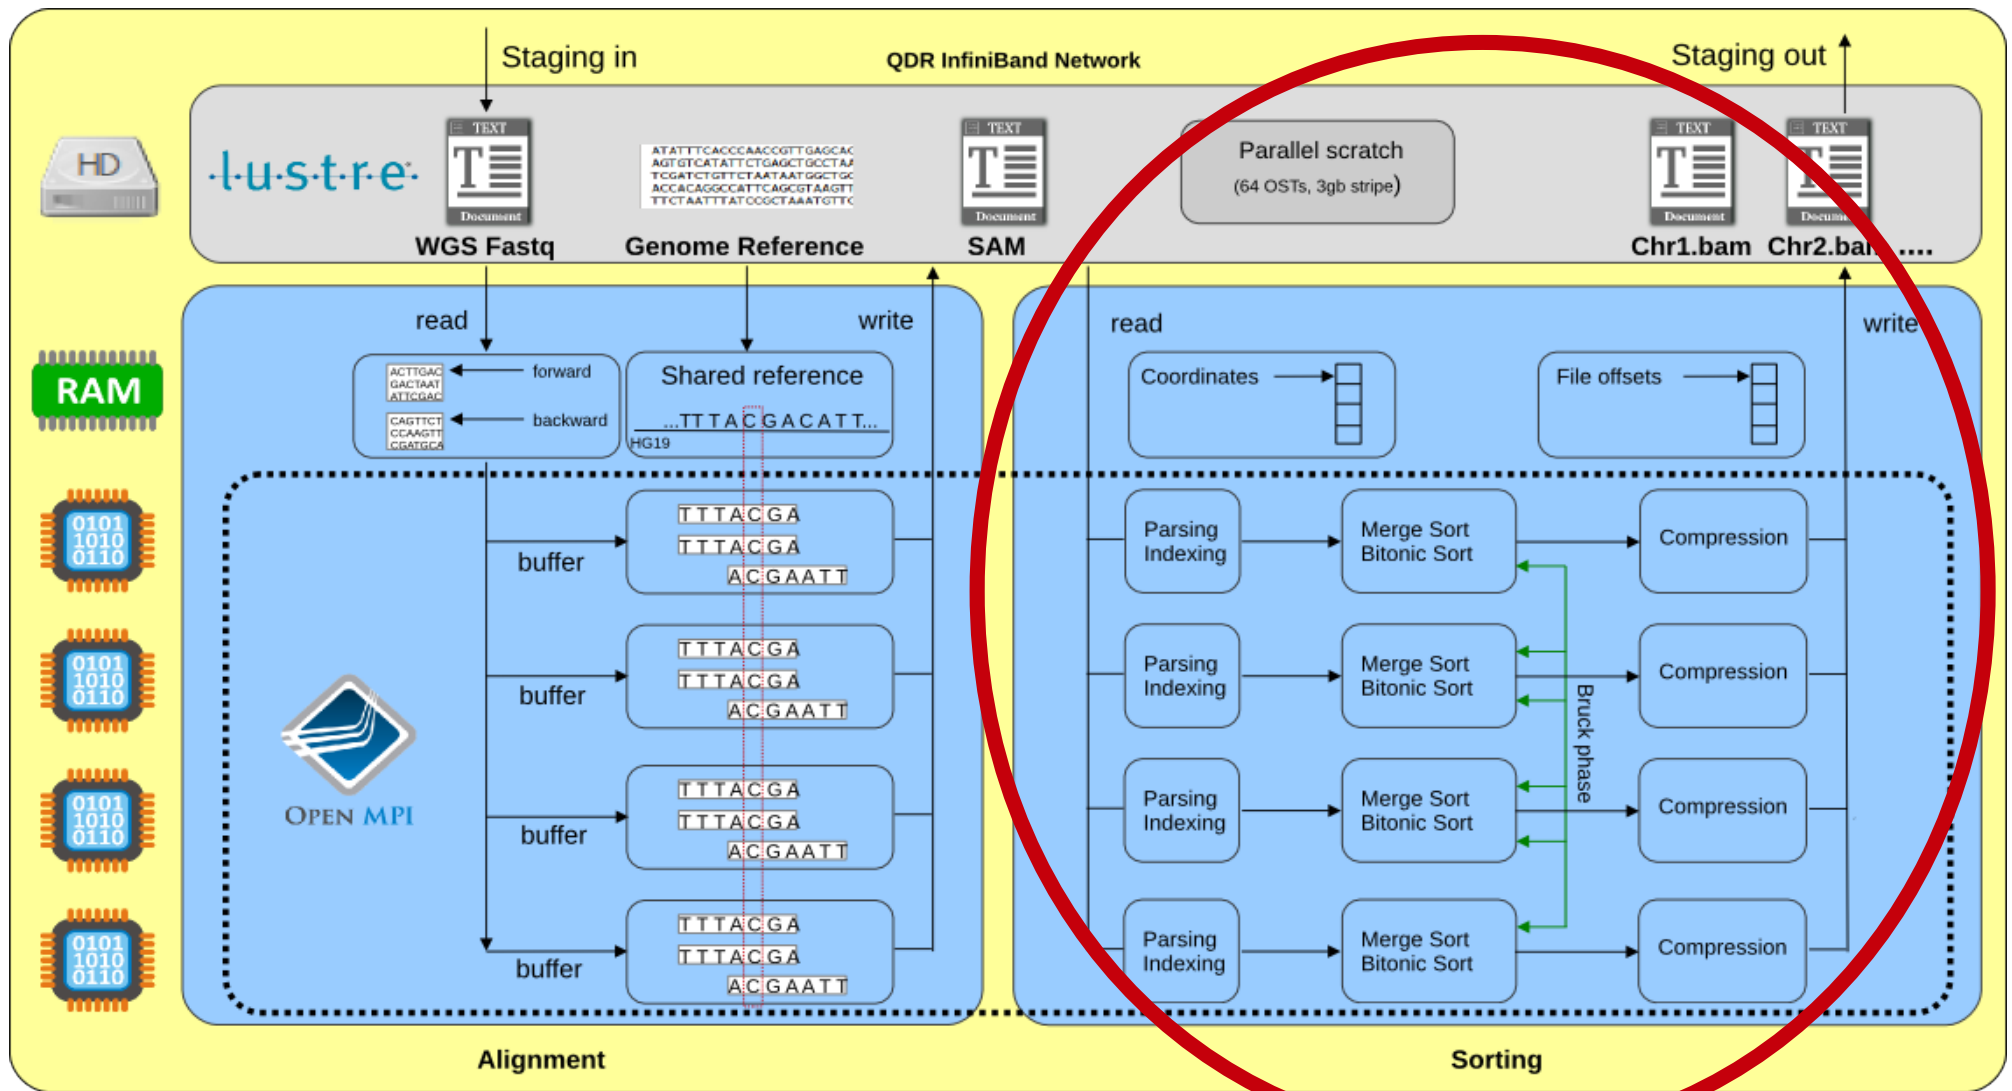


Now the alignment is CPU bound

Memory impact < 2gb per jobs

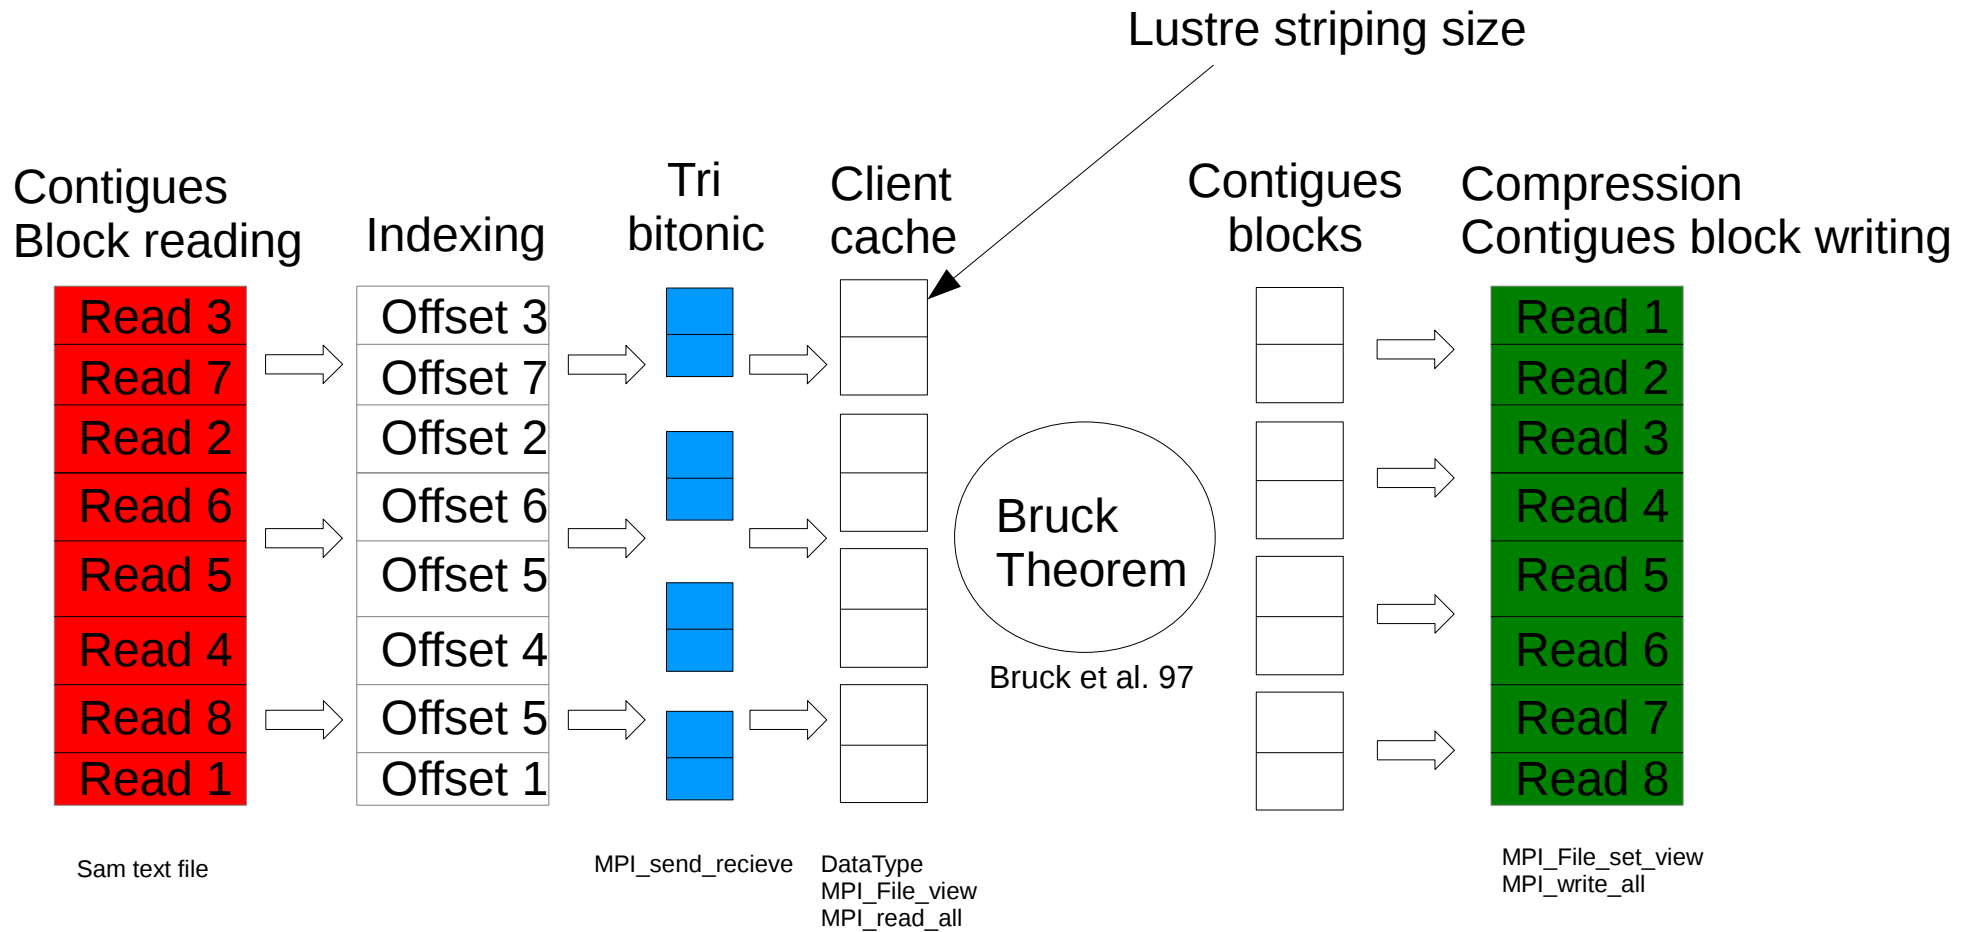
• Sorting





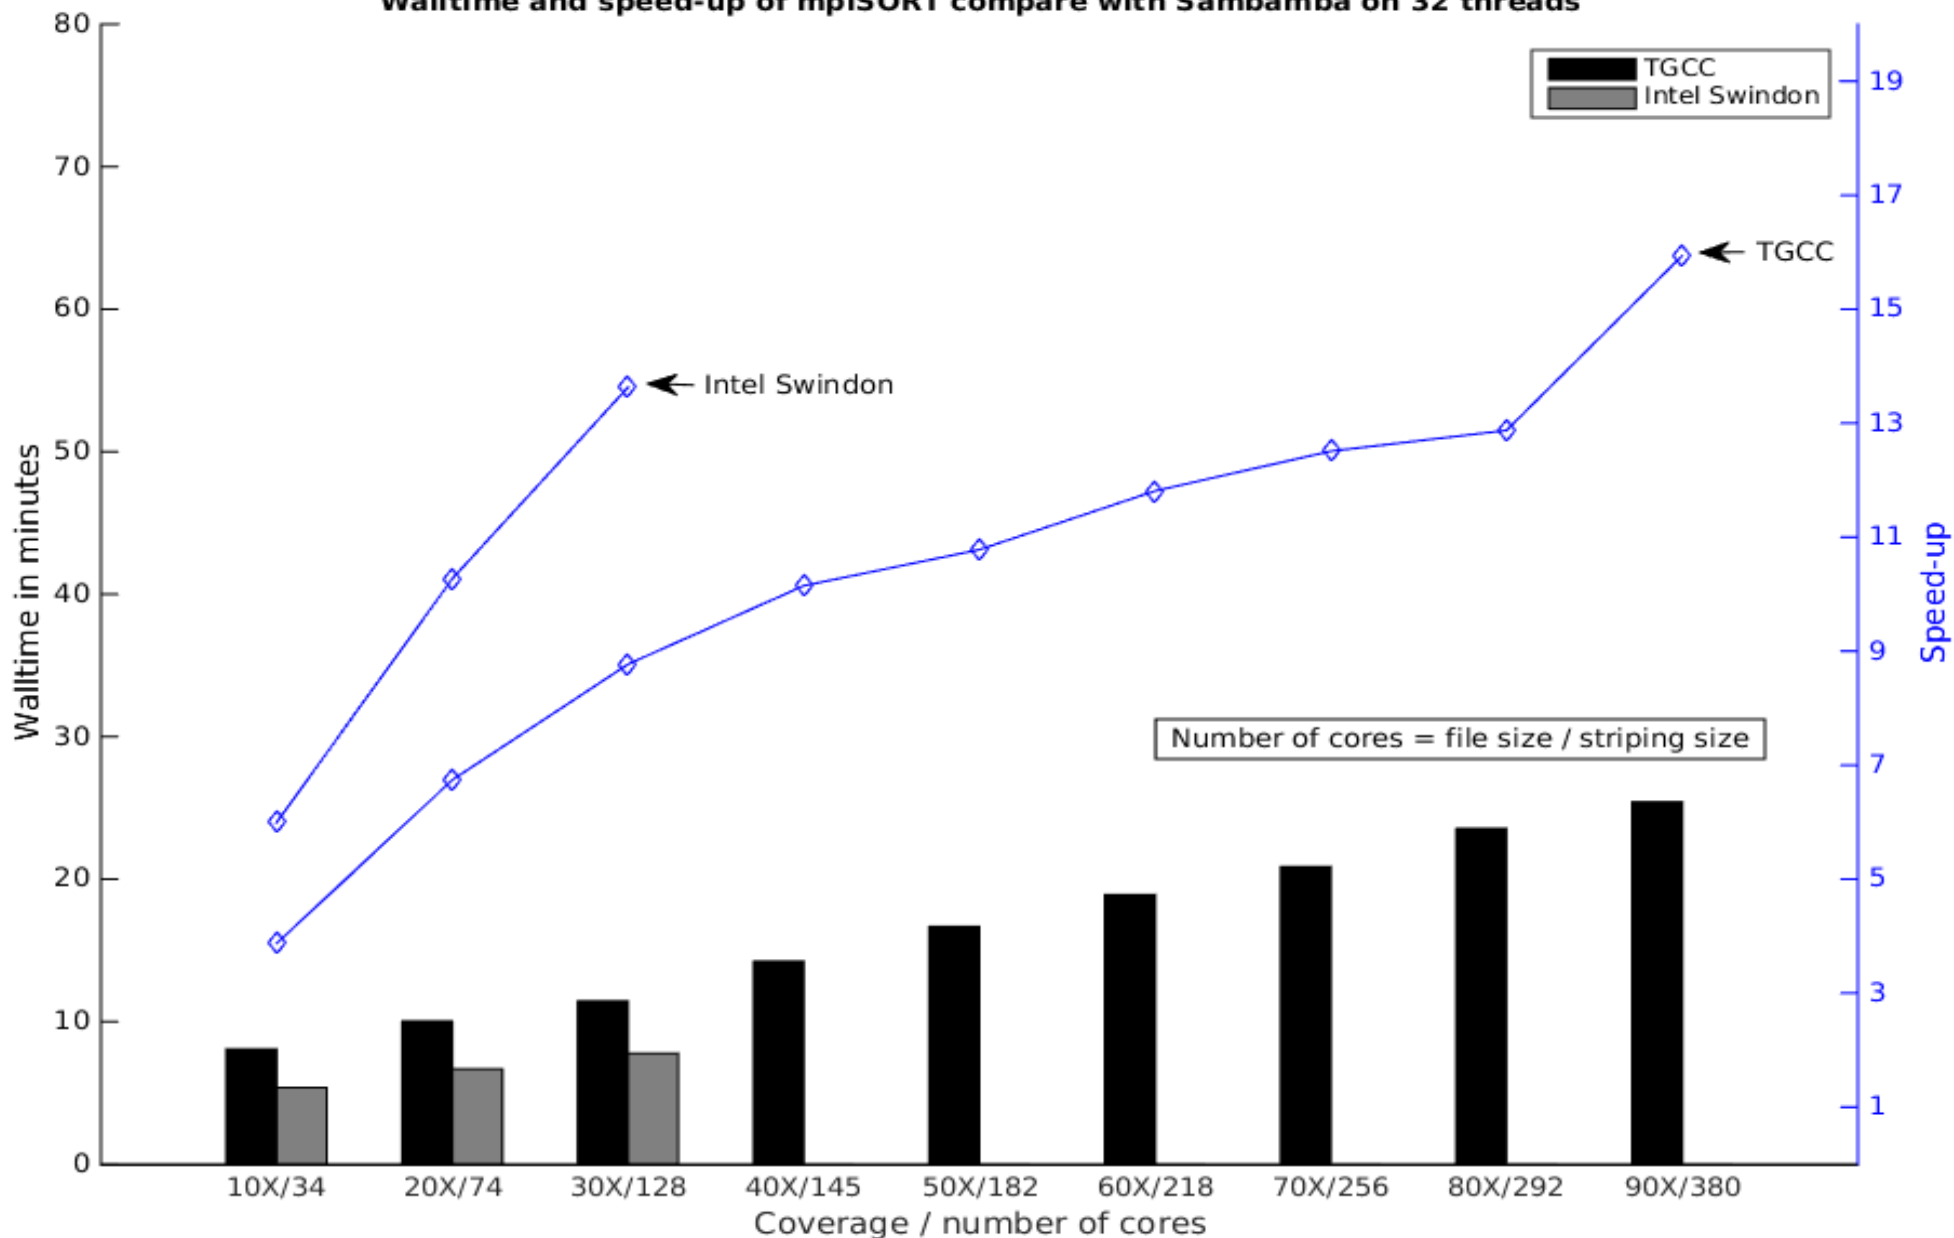
<https://github.com/InstitutCurie/QUASART>

- Bitonic sorting



Repeated reads
local data

Walltime and speed-up of mpiSORT compare with Sambamba on 32 threads



25 minutes for sorting a 120X (1,2 TB) with 512 cpu, RAM per CPU < 4Gb

32 threads

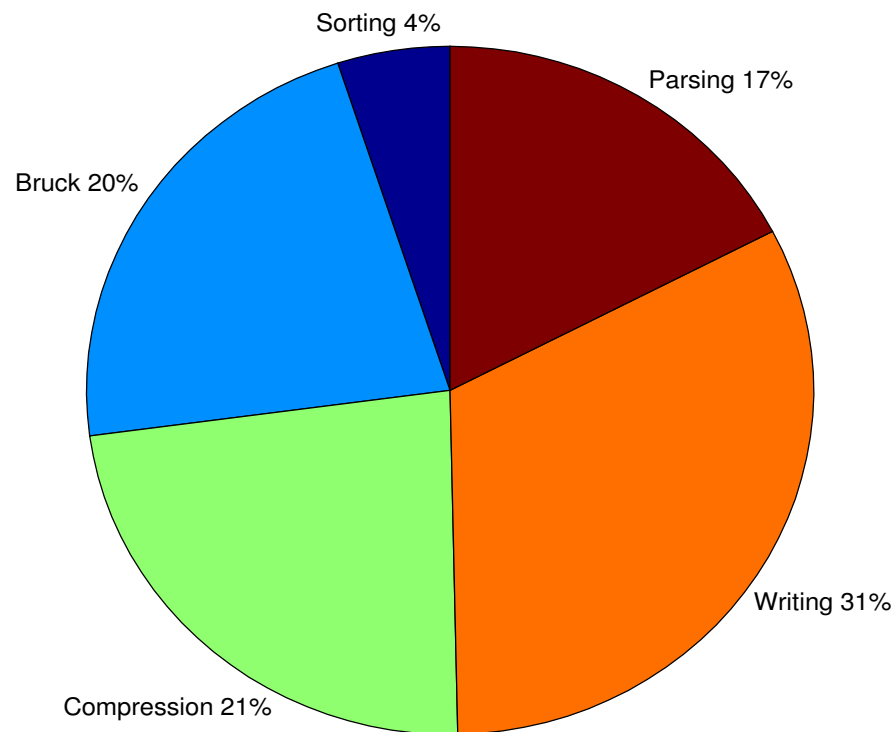


Wall time * (nb cores)



Coverage	Sambamba (s)	mpiSort (s)	CPUs	speed-up	efficiency (%)
10X	1875	484	36	3.87	70.96
20X	4058	602	74	6.74	65.69
30X	6030	687	128	8.77	54.42
40X	8166	854	145	9.56	52.61
50X	10222	1000	182	10.22	44.36
60X	12590	1134	218	11.1	38.63
70X	14452	1246	256	11.59	30.69
80X	16625	1341	292	12.39	22.44
90X	24860	1574	380	15.79	14.94
120X	fail	1836	512	x	x

Analyse of Sorting algorithm for 93X WG



- **Configuration**

Lustre configuration

```
max_cached_mb: 48301  
max_read_ahead_mb=40  
max_pages_per_rpc=256  
max_rpcs_in_flight=32
```

File striping

```
lfs setstripe -s 3g -c 143 .
```

MPI tuning info

```
MPI_Info_create(&finfo);  
MPI_Info_set(finfo,"striping_factor","143");  
MPI_Info_set(finfo,"striping_unit","1073741824");
```

```
MPI_Info_set(finfo,"access_style","read_mostly");  
MPI_Info_set(finfo,"nb_proc","143");  
MPI_Info_set(finfo,"cb_block_size","1073741824");  
MPI_Info_set(finfo,"cb_buffer_size","1073741824");
```

- **Conclusion**

- 2 major bottlenecks hamper the scaling of NGS analysis : the IO and memory
- Lustre has the ability to efficiently share files across the network enabling the alignment to scale thousand of cores
- The striping tuning of Lustre give me the ability to align the chunks of the file to sort and improve the performance
- The low latency fabric enable new parallel algorithms (Bitonic Sort) implementation

- **Next step**

- Challenge new architectures (Intel Broadwell)

• Thanks

– From Institut Pasteur :

- Nicolas Joly



– From TGCC :

- Claude Scarpelli



– From Intel

- Gabriele Paciucci, Kristina M Kermanshahche



– From Institut Curie :

- Philippe Hupé (Supervisor), François Prud'Homme (IT), Maxime Chevilliot(IT)

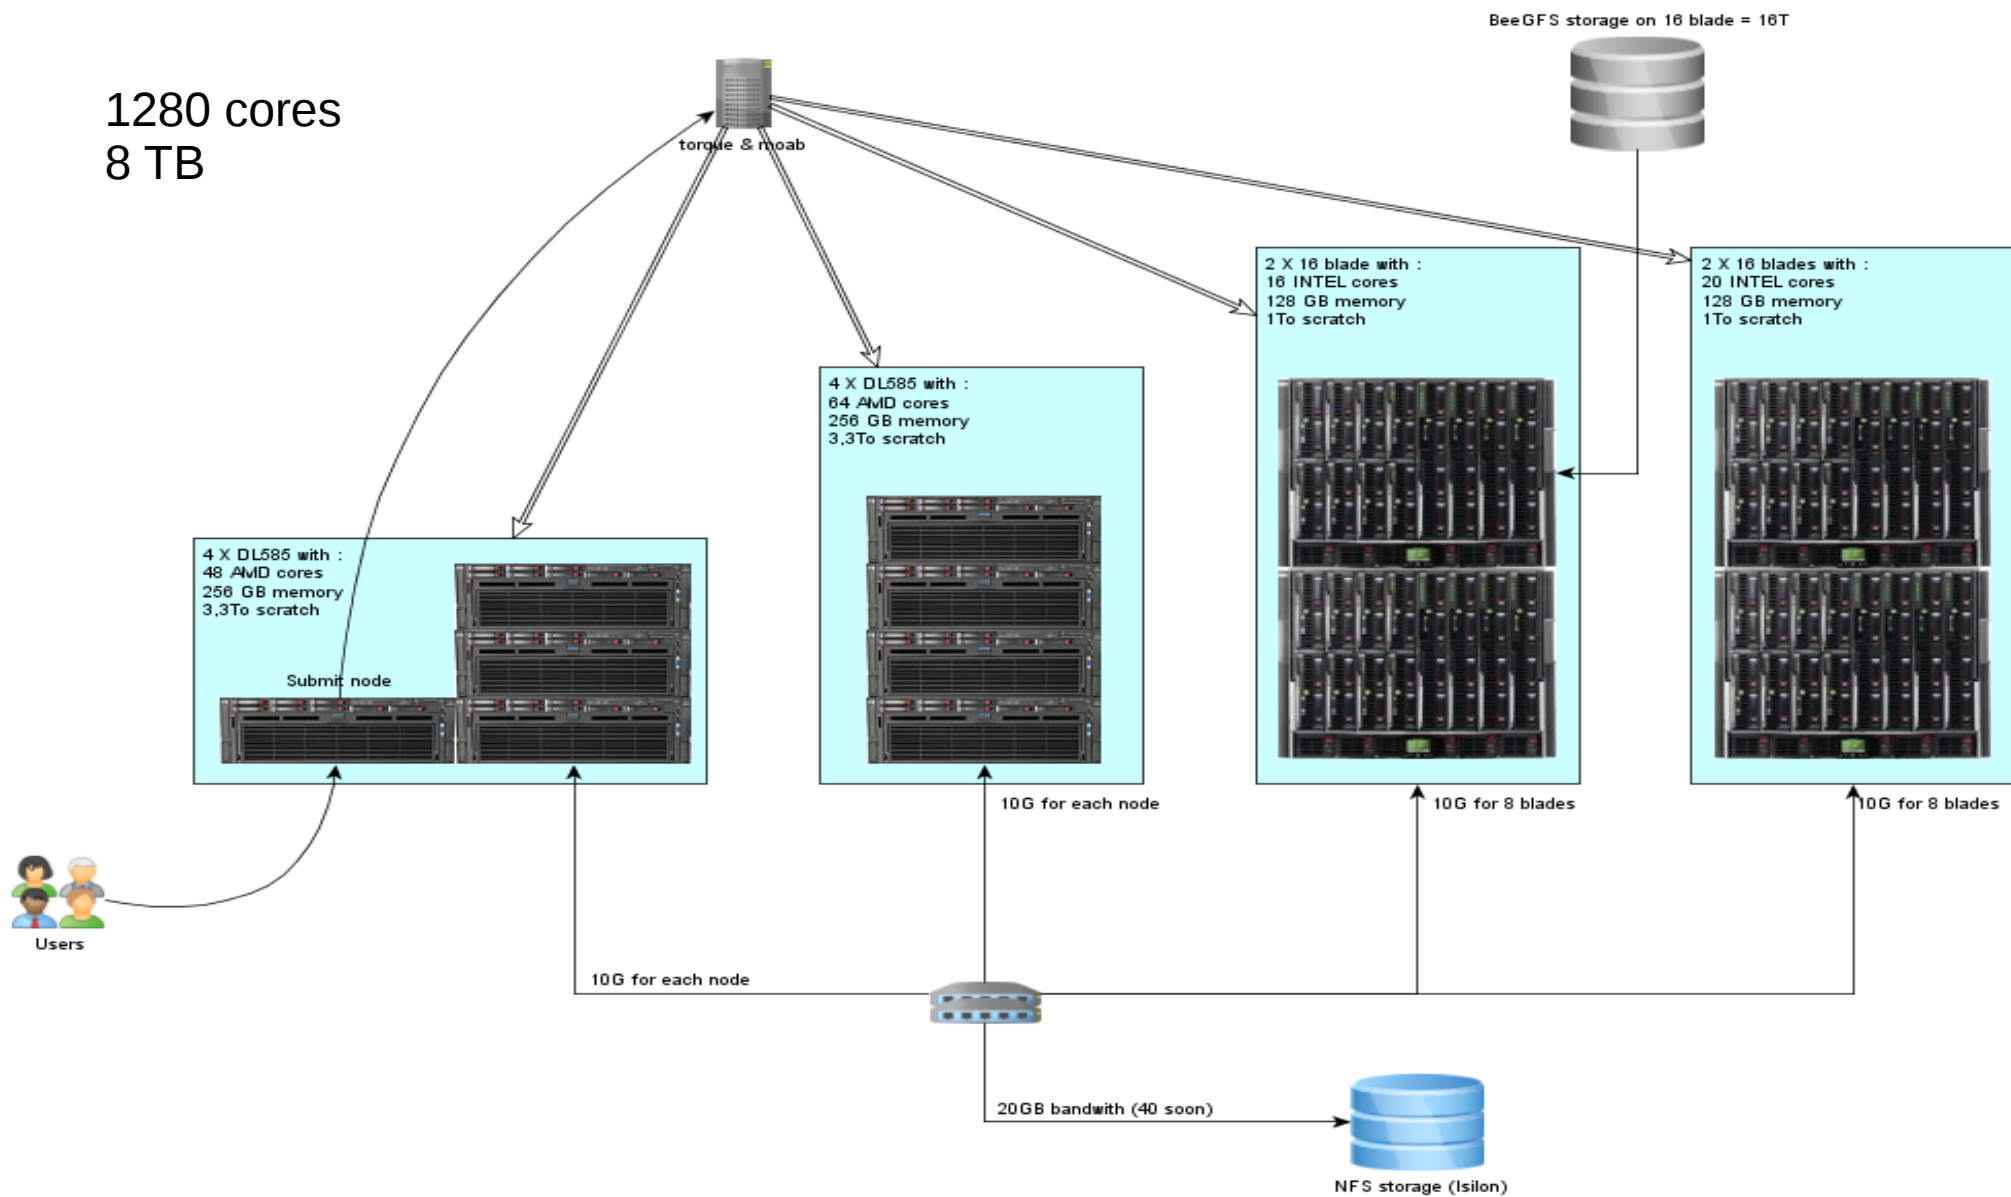
– From University of Paris Descartes :

- Nicolas Fedy, Leonor Sirotti, Thomas Magalhaes, Paul Paganiban



• Curie architecture

1280 cores
8 TB



• Exemple of Intel clinical architecture

