



# **The GeneROOT Project Status and Plans**

Fons Rademakers

CERN openlab Chief Research Officer



# GeneROOT - Using ROOT for Handling Genomics Data


# King College London - TwinsUK Project

- Collaboration between the KCL and CERN openlab
- Try to optimize genomics data storage and processing using HEP tools
- Working on a local 400TB copy of TwinsUK data
  - 750 Monozygotic twins
  - 900 Dizygotic twins
  - 138 Singletons




# Rapidly Increasing Amount of Genomics Data

- Next generation Sequencing (NGS)
  - Dramatic increase in the amount of data
  - Improved data confidence
- NGS is enabler for more sophisticated research questions in Genomics



**Issue: Leaps in sequencing technology have outperformed advances in computing**

### Growth of DNA Sequencing



Growth of DNA sequencing data both in terms sequenced human genomes and total sequencing capacity

# SAM - Sequence Alignment/Map

- Text-based format for storing biological sequences aligned to a reference one

- SAM - Plaintext
- BAM - Binary Compressed



Challenge: data size is extremely large (about 500GB for a single human genome)

# SAM Example

```
@SQ      SN:chrM LN:16571
@SQ      SN:chrX LN:155270560
SOLEXA-1GA-2_2  0   chr1  10145  25  36M  *  0  0  AACCCCTAACCCCTAACCCCTAACCCCTA  hhhhHcWhhHTghcKA_ONhAAEEBZ
SOLEXA-1GA-2_2  0   chr1  10148  25  36M  *  0  0  CCCTAACCCCTAACCCCTAACCCCTAACC  hbfhhhXUYhT_ULZdLRTKNIMIKG  NM:i:0
SOLEXA-1GA-2_2  16  chr1  10149  25  36M  *  0  0  CCAAACACTAACCCCTAACCCCTAACCC  ><>B@>?>?D>>?B?D>DBC?E@BDH  NM:i:1  X1:i:1
SOLEXA-1GA-2_2  0   chr1  10150  25  36M  *  0  0  CTAACCCTAACCCCTAACCCCTAACCCCT  hhW_X]MXNOHQQWMILHGIFMJGJ
```

The file consists of a header section (lines starting with a @)  
and an alignment section with 11 mandatory fields + several optional fields

A human genome SAM file consist of about 1 billion lines in the alignment section and is about 500GB large

# BAM - Binary Alignment/Map

- BGZF is block compression implemented on top of the standard gzip file format
- File is gunzip compatible
- Each 64KB block of data is compressed and added to the file
- Using an .bai index file, random access is supported in the BAM file

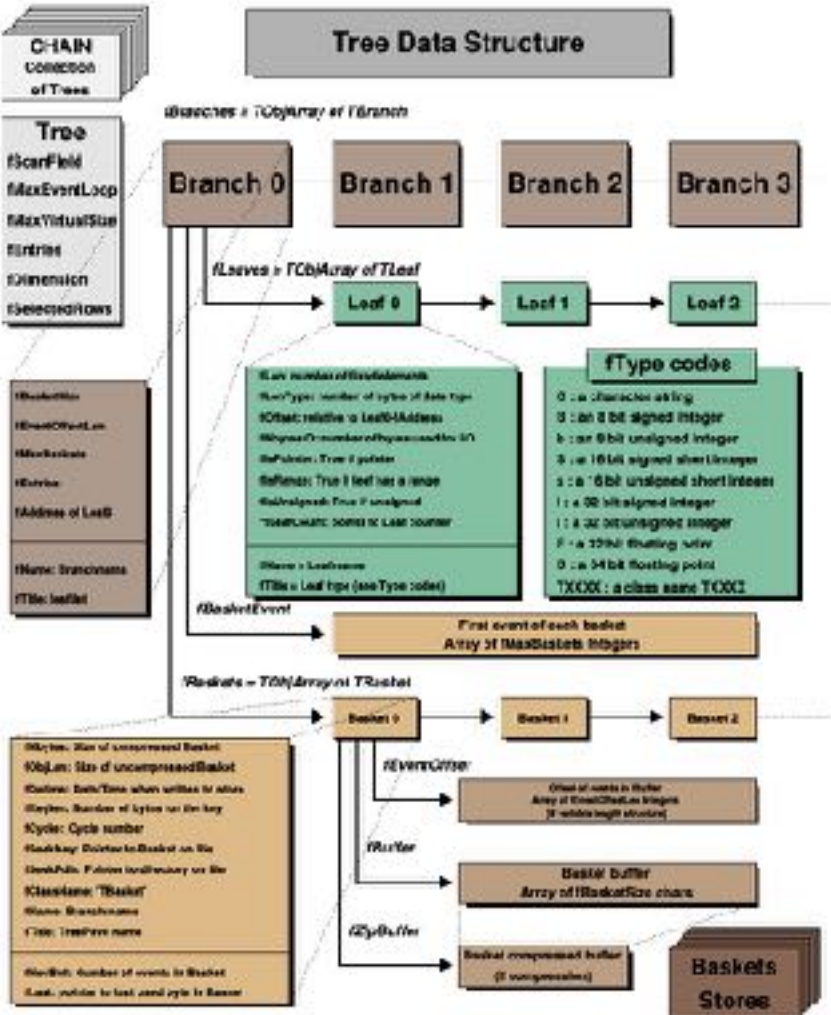


# ROOT Framework

Don't reinvent the wheel



Most of the challenges with massive data processing are common to **HEP**. **ROOT** has decades of experience in software design and optimisation



# Conversion - SAM to ROOT

We convert from TAB-delimited SAM format to a ROOT TTree

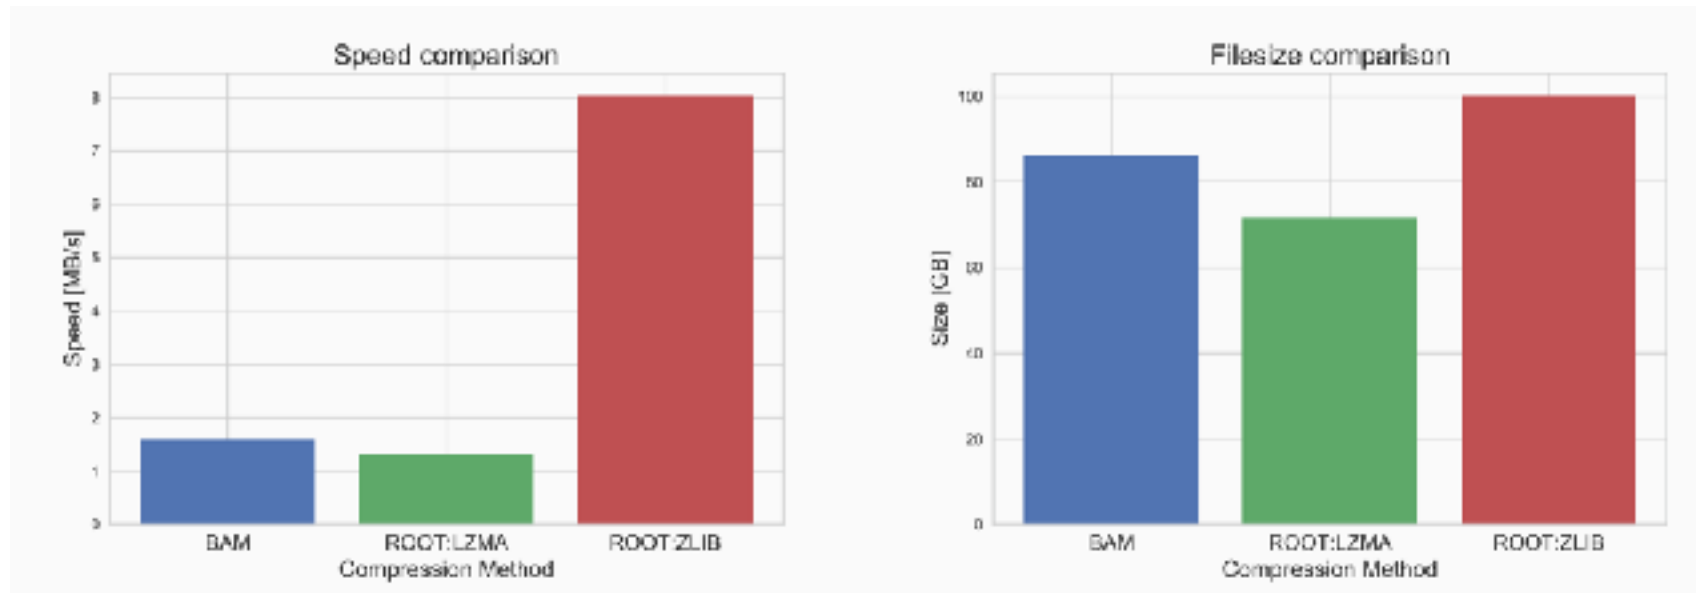
```
SOLEXA-1GA-2_2_FC28EMB:5:251:979:328 0 chr1 18145 25 36M * 0 0 AACCCCTAACCCCTAACCCCTAACCCCTAACCTAACCT hhhhHcwhhHTghcKA_OnhAAEEBZE?H?CBC?DA NM:i:1 X1:i:1 MD:Z:33A
SOLEXA-1GA-2_2_FC28EMB:5:182:214:278 0 chr1 18148 25 36M * 0 0 CCTAACCCCTAACCCCTAACCCCTAACCCCTAACCTAAC hbFhhXUJYhT_ULZdLRtNIMIKKJCHFFJQJN NM:i:0 X0:i:1 MD:Z:36
SOLEXA-1GA-2_2_FC28EMB:5:195:284:685 16 chr1 18149 25 36M * 0 0 CCAAACCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC >=>Bg>?>?d>>?B?d>DBC?EgBDIAKCEKERLOO NM:i:1 X1:i:1 MD:Z:29A
SOLEXA-1GA-2_2_FC28EMB:5:35:583:827 0 chr1 18150 25 36M * 0 0 CTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCA NHM_X]M9N0HQ0MILHGTFMDGJLCFGGJAKIEH NM:i:1 X1:i:1 MD:Z:18A
SOLEXA-1GA-2_2_FC28EMB:5:248:138:724 0 chr1 18152 25 36M * 0 0 AACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCTA hchPhc__cWS [VR0bRXDT0USJLXDAGLGFMCg NM:i:1 X1:i:1 MD:Z:17C
SOLEXA-1GA-2_2_FC28EMB:5:236:644:187 16 chr1 18154 25 36M * 0 0 CCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAAC ICF=Ag8HGEGGDFLTKYYGD^CalMaShfNhhhh NM:i:0 X0:i:1 MD:Z:36
SOLEXA-1GA-2_2_FC28EMB:5:165:628:78 16 chr1 18155 25 36M * 0 0 CCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC >D?BAEAA?E=JGGBK>FKDGFJPVWISFTT\Q0Mch NM:i:1 X1:i:1 MD:Z:11A
SOLEXA-1GA-2_2_FC28EMB:5:188:485:455 16 chr1 18156 25 36M * 0 0 CTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC CFVUIIIIbGPROGRhRhhThhhhhhhhhhhhhfhhh NM:i:1 X1:i:1 MD:Z:12A
SOLEXA-1GA-2_2_FC28EMB:5:248:581:237 0 chr1 18158 25 36M * 0 0 AACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCATA hhhchg_ORNbx]RMZLREQWTFNIFLDPHDDDEDKL NM:i:1 X1:i:1 MD:Z:33A
SOLEXA-1GA-2_2_FC28EMB:5:258:882:389 16 chr1 18215 25 36M * 0 0 CTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACAA IMZD>?EGORRhhDQWRZKhhhhhhhhhhhhhh NM:i:1 X1:i:1 MD:Z:25A
SOLEXA-1GA-2_2_FC28EMB:5:197:589:878 16 chr1 18216 25 36M * 0 0 TAAACCCGAAACCCGAAACCCCTAACCCCTAACCCCTAACAA <IJELCgSCUMY?R?D\UPW^gd]h00N\hhhhhhLh NM:i:1 X1:i:1 MD:Z:23G
SOLEXA-1GA-2_2_FC28EMB:5:168:888:612 0 chr1 18217 25 36M * 0 0 AACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACAA hNNhV\0 [MUNYQKOWNSIIDNFOPIJHIAI@GDJG NM:i:1 X1:i:1 MD:Z:15T
SOLEXA-1GA-2_2_FC28EMB:5:249:922:888 16 chr1 18217 25 36M * 0 0 AACCTAACCCCTAACCCCTAACCCCTAACCCCTAACAA h0iQRBhhB` `Ah` hhhhhhhhhhhhhhhhhhhhh NM:i:1 X1:i:1 MD:Z:8A2
SOLEXA-1GA-2_2_FC28EMB:5:13:922:731 0 chr1 18236 25 36M * 0 0 GACCCAAACCCCTAACCCCTAACCCCTAACCCCTAACCCTAACC NHWZaPTSZHfUMMS0JTJELKELKXDJ>JPCETI NM:i:1 X1:i:1 MD:Z:8G4
SOLEXA-1GA-2_2_FC28EMB:5:62:877:892 16 chr1 18237 25 36M * 0 0 ACCCCAAACCCCTAACCCCTAACCCCTAACCCCTAACCCTAACCC CHLKEQGDUCQ`GYdhF [hXhhHhhhhhhhhhh NM:i:1 X1:i:1 MD:Z:16C
SOLEXA-1GA-2_2_FC28EMB:5:172:417:550 16 chr1 18242 25 36M * 0 0 AACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCTAACCC B@BA=>EG>B@CANEDBEKMGROMVJTORSNKh NM:i:0 X0:i:1 MD:Z:36
SOLEXA-1GA-2_2_FC28EMB:5:129:615:872 0 chr1 18243 25 36M * 0 0 ACCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC ghhhhhZ`h]hieh\^SZ_QOUVL\IQwLTQJOMII NM:i:0 X0:i:1 MD:Z:36
SOLEXA-1GA-2_2_FC28EMB:5:223:388:145 0 chr1 18244 25 36M * 0 0 CCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCCCT g [WNNHUM0dWTIRNJSTIGHILLJJBFB>?B?F NM:i:0 X0:i:1 MD:Z:36
SOLEXA-1GA-2_2_FC28EMB:5:196:958:794 16 chr1 18247 25 36M * 0 0 AAAACCCAAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAAC LV_KJgc~Pf00ghThKahh` hhhhhhhhhhhhh NM:i:1 X1:i:1 MD:Z:28A
SOLEXA-1GA-2_2_FC28EMB:5:6:541:676 16 chr1 18316 25 36M * 0 0 ATCACC AACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC P?XLX?UGCKEhQeU [hdWNa [hYqGhhhhVhIhhh NM:i:1 X1:i:1 MD:Z:32A
SOLEXA-1GA-2_2_FC28EMB:5:32:129:784 0 chr1 18316 25 36M * 0 0 CACCCAAACCCCTAACCCCTAACCCCTAACCCCTAACCC hhhhhhhhhhhcchhhhhhhhhZh]lhf^ghQ^THY NM:i:1 X1:i:1 MD:Z:8C2
```

ROOT handles compression, memory buffers, datatypes, endian-ness, etc.

The columns are defined using a C++ class.

# Performance - SAM to ROOT

There is a tradeoff between compression and read/write speed for 100GB file.



With ZLIB compression -> 4 times faster

With LZMA compression conversion -> 15% smaller

# View - Random Access

- We also need to be able to view the information as fast as possible
- ROOT columnar structure allow us to just look at the **chromosome** and **position** columns to optimize performance

SOLEXA-1GA-2_2_FC28EMB;5:35:583:827	0	chr1	18158	25	36N	*	0	0	CTAACCCCTAACCTAACCCCTAACCCCTAACCCCTAACCC	HFk_XfP0wQhQqWMLHGIFPUGJLCPGGJAKIEH	NN:i:1	X1:i:1	ND:Z:18A
SOLEXA-1GA-2_2_FC28EMB;5:248:138:724	0	chr1	18152	25	36N	*	0	0	AACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC	hchPhc___cW5 [VR06-RX3T0USJLXDA@LGFN0@	NN:i:1	X1:i:1	ND:Z:17C
SOLEXA-1GA-2_2_FC28EMB;5:236:644:187	36	chr1	18154	25	36N	*	0	0	CCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC	ICP-A6BHG6G9FLIKKYGD~CaMhSaSHfMhhh	NN:i:0	X8:i:1	ND:Z:36
SOLEXA-1GA-2_2_FC28EMB;5:166:628:78	36	chr1	18155	25	36N	*	0	0	CCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC	>078AEAA7E~JGGB9<FXDGFJPHV5FTT\QONch	NN:i:1	X1:i:1	ND:Z:11A
SOLEXA-1GA-2_2_FC28EMB;5:188:485:455	36	chr1	18156	25	36N	*	0	0	CTAACCCCTAACCCCTAACCCCTAACCCCTAACCC	CFVIIIIbGPR0GhRhhIthhhhhhhhhhhfhhh	NN:i:1	X1:i:1	ND:Z:12A
SOLEXA-1GA-2_2_FC28EMB;5:240:581:237	0	chr1	18158	25	36N	*	0	0	AAACCCCTAACCCCTAACCCCTAACCCCTAACCCATA	Hhchg_0RNlx]RMZLRE@NNTMFL0PMDD0EdKL	NN:i:1	X1:i:1	ND:Z:33A
SOLEXA-1GA-2_2_FC28EMB;5:258:802:389	16	chr1	18215	25	36N	*	0	0	CTAACCCCTAACCCCTAACCCCTAACCCCTAAC	IMZ0>?EGDRRhh0XVNRZK-Hhhhhhhhhhhhh	NN:i:1	X1:i:1	ND:Z:25A
SOLEXA-1GA-2_2_FC28EMB;5:197:589:878	36	chr1	18216	25	36N	*	0	0	TAAACCGAAACCGAAACCCCTAACCCCTAACCCCTAAC	<EJELCgSQUY?R7DVUPV~gdIH0XV\HhhhhLh	NN:i:1	X1:i:1	ND:Z:236
SOLEXA-1GA-2_2_FC28EMB;5:160:888:632	0	chr1	18217	25	36N	*	0	0	AACCCCTAACCCCTAACCCCTAACCCCTAACCCATAACA	H-Hh\0 [MUNYQKQMSIIDNFOPJHIAI@gD3G	NN:i:1	X1:i:1	ND:Z:15T
SOLEXA-1GA-2_2_FC28EMB;5:240:922:808	36	chr1	18217	25	36N	*	0	0	AACACTAACCCCTAACCCCTAACCCCTAACCCCTAACCA	h0wQRhB" `Ah" hhhhhhhhhhhhhhhhhhhhh	NN:i:1	X1:i:1	ND:Z:8A2
SOLEXA-1GA-2_2_FC28EMB;5:13:922:731	0	chr1	18236	25	36N	*	0	0	GACCCAAACCCCTAACCCCTAACCCCTAACCCCTAAC	NHWZaPTSZHTUMSDJTJELKJELKDJ>3PCEI	NN:i:1	X1:i:1	ND:Z:854
SOLEXA-1GA-2_2_FC28EMB;5:62:877:892	36	chr1	18237	25	36N	*	0	0	ACCCCAAACCCCTAACCCCTAACCCCTAACCCCTAACCC	CHLkEGGdUcCq" GyAhF [n0nhhhhhhhhhhh	NN:i:1	X1:i:1	ND:Z:16C
SOLEXA-1GA-2_2_FC28EMB;5:172:417:558	36	chr1	18242	25	36N	*	0	0	AAACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAAC	BgRA>=EG>Bg@NBEDBKBNCR0M0VJTORSHKh	NN:i:0	X8:i:1	ND:Z:36
SOLEXA-1GA-2_2_FC28EMB;5:129:615:872	0	chr1	18243	25	36N	*	0	0	ACCCCTAACCCCTAACCCCTAACCCCTAACCCCTAACCC	ghhhhhZ^hI hwhV^*S_Z_Q0uLV\IQWLTQJQNI	NN:i:0	X8:i:1	ND:Z:36

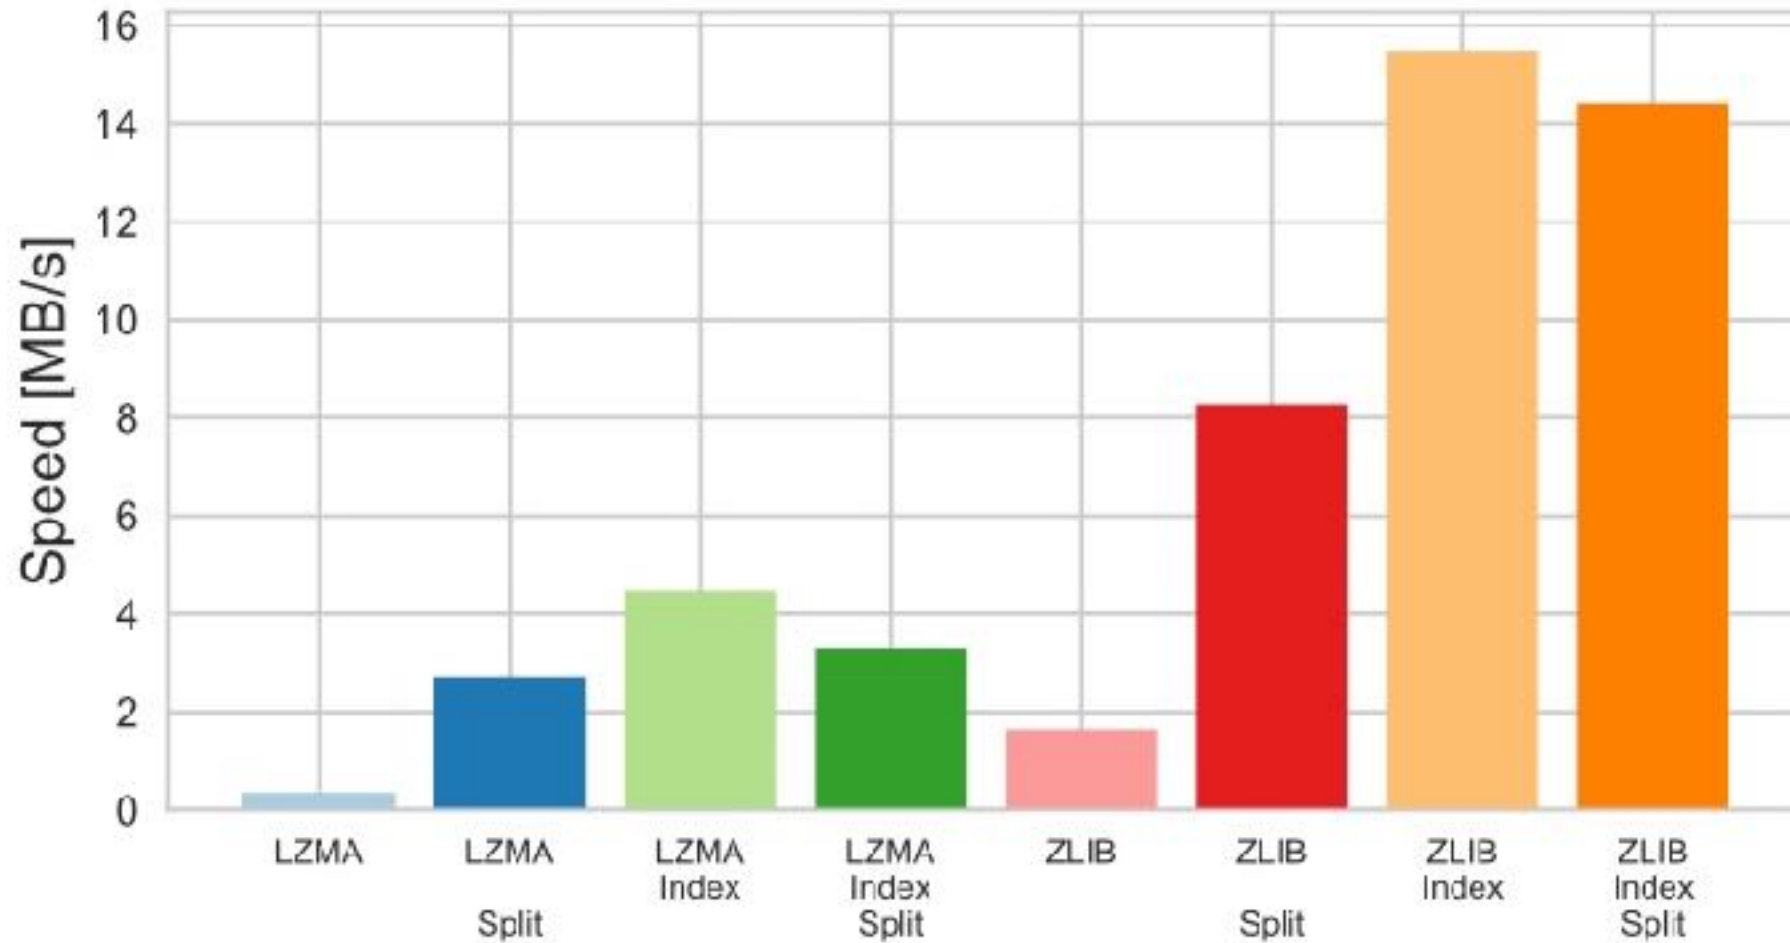
# Indexing For Fast Access

- Improve read speed by using an index
  - For random access BAM format needs a BAI index file
  - For ROOT implemented a RAMIndex, sparse index in combination with fast columnar ROOT file scanning
  - 16 bytes per entry, compressed

```
class RAMIndex {  
private:  
    typedef std::pair<int,int> Key_t;           // refid (of rname) and pos  
    typedef std::map<Key_t,Long64_t> Index_t; // map of Key_t and TTree entry number  
  
    Index_t fIndex;  
    ...  
};
```



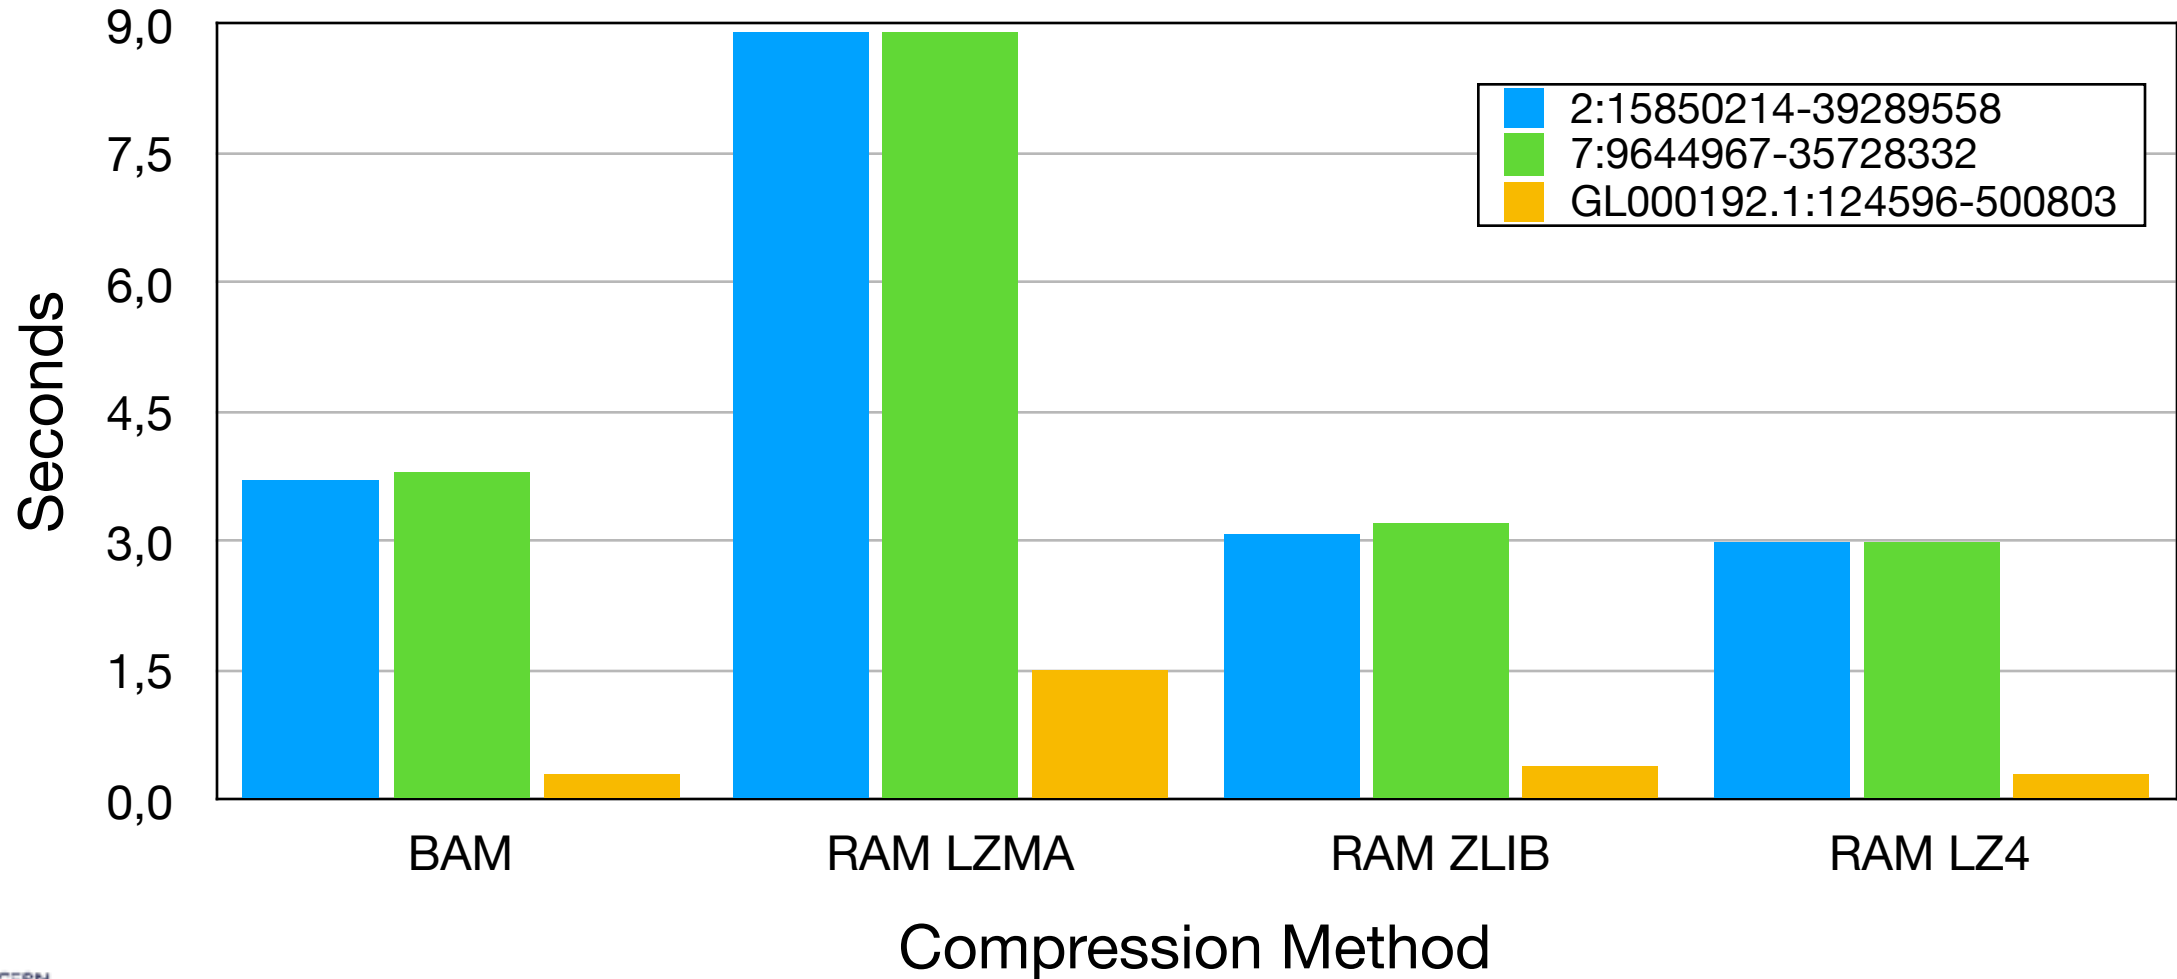
# View - Performance



Median read speed for ramtools across different parameter settings

# View - Performance

Viewing 3 different regions from a 9-18 GB file



# Future Work

- Increase file sample - The study was done on only one base SAM file (subsamped to 10% and 1%). Conversion, indexing and views should be run on different SAM files so sample bias is reduced
- Extended format comparison - The current work compared only to the well established BAM format. However, comparisons to formats such as CRAM should be made to see the relative performances in compression rate and speed and read speed, respectively
- Support for additional operations - Add to ramtools support for sort, merge, split or stats.
- FASTQ/FASTA conversion - The raw sequencer data format. Simpler but closely resembling SAM. Most of the TTree advantages apply also to this type of data. In fact, recent research in formats like LFQC has many analogies to how ROOT branches are used to compress data fields separately, optimizing compression and read speed.



# CERN openlab GeneROOT Technology Transfer Benefits

- Additional use case for CERN's ROOT technology
- Return flow of know-how benefiting the ROOT User community
- CERN openlab provides a doctoral student who gains valuable experience
- Entire Omics community would benefit from improved analysis tools to handle rapidly growing amounts of data
- Project is joint effort between CERN openlab, CERN medical applications group, King's College London
  
- Big thank you to my 2017 openlab summer student  
Jose Javier Gonzalez Ortiz