










TECHNICAL NOTE

HTSlib: C library for reading/writing high-throughput sequencing data

James K. Bonfield ^{1,†}, John Marshall ^{2,†}, Petr Danecek ^{1,†}, Heng Li ^{3,4}, Valeriu Ohan ¹, Andrew Whitwham ^{1,*}, Thomas Keane ⁵ and Robert M. Davies ¹

¹Wellcome Sanger Institute, Wellcome Genome Campus, Hinxton, Cambridgeshire CB10 1SA, UK; ²Wolfson Wohl Cancer Research Centre, Institute of Cancer Sciences, University of Glasgow, Switchback Road, Glasgow, G61 1QH, UK; ³Department of Data Sciences, Dana-Farber Cancer Institute, 450 Brookline Avenue, Boston, MA 02215, USA; ⁴Department of Biomedical Informatics, Harvard Medical School, 10 Shattuck Street, Boston, MA 02215, USA and ⁵EMBL-EBI, Wellcome Genome Campus, Hinxton, Cambridgeshire, CB10 1SD, UK

*Correspondence address. Andrew Whitwham, Wellcome Sanger Institute, Wellcome Genome Campus, Hinxton, Cambridgeshire, CB10 1SA, UK. Tel: +44 (0)1223 834244; E-mail: samtools@sanger.ac.uk  <http://orcid.org/0000-0001-8117-400X>

[†]Contributed equally.

Abstract

Background: Since the original publication of the VCF and SAM formats, an explosion of software tools have been created to process these data files. To facilitate this a library was produced out of the original SAMtools implementation, with a focus on performance and robustness. The file formats themselves have become international standards under the jurisdiction of the Global Alliance for Genomics and Health. **Findings:** We present a software library for providing programmatic access to sequencing alignment and variant formats. It was born out of the widely used SAMtools and BCFTools applications. Considerable improvements have been made to the original code plus many new features including newer access protocols, the addition of the CRAM file format, better indexing and iterators, and better use of threading. **Conclusion:** Since the original Samtools release, performance has been considerably improved, with a BAM read-write loop running 5 times faster and BAM to SAM conversion 13 times faster (both using 16 threads, compared to Samtools 0.1.19). Widespread adoption has seen HTSlib downloaded >1 million times from GitHub and conda. The C library has been used directly by an estimated 900 GitHub projects and has been incorporated into Perl, Python, Rust, and R, significantly expanding the number of uses via other languages. HTSlib is open source and is freely available from htslib.org under MIT/BSD license. **Keywords:** samtools, bcftools, high-throughput sequencing, next generation sequencing, variant calling, data analysis.

Background

When the 1000 Genomes Project [1] was launched in early 2008, there were many short-read aligners and variant callers. Each of them had its own input or output format for limited use cases, and interoperability was a major challenge. Users were forced to implement bespoke format converters to bridge tools, and because formats encoded different information, this was time consuming, laborious, and sometimes impossible. This fragmented

ecosystem hampered the collaboration between the participants of the project and delayed the development of advanced data analysis algorithms.

In a conference call on 21 October 2008, the 1000 Genomes Project analysis subgroup decided to take on the issue by unifying a variety of short-read alignment formats into the Sequence Alignment/Map format (SAM). Towards the end of 2008 the subgroup announced the first SAM specification, detailing a text-

Received: 16 December 2020; Revised: 18 January 2021; Accepted: 28 January 2021

© The Author(s) 2021. Published by Oxford University Press GigaScience. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

based SAM format and its binary representation, the BAM format [2]. SAM/BAM quickly replaced all the other short-read alignment formats and became the *de facto* standard in the analysis of high-throughput sequence data. In 2010, a Variant Call Format (VCF) was introduced for storing genetic variation [3]. Later, in 2011, as the number of sequenced samples increased and the text format proved too slow to parse, a binary version BCF (see Supplementary Fig. S1) was developed [4].

The SAM/BAM format originally came with a reference implementation, SAMtools [2, 5], and VCF/BCF with VCFtools [3] and BCFtools (then part of the SAMtools package) [6]. Numerous other tools have been developed since then in a wide variety of programming languages. For example, HTSJDK is the Java equivalent [7] and is used extensively in Java applications; Sambamba [8] is written in the D language and focuses primarily on efficient multi-threaded work; Scramble [9] has BAM and SAM capability and is the primary source for experimental CRAM [10, 11] development; and JBrowse [12] implements read-only support for multiple formats in JavaScript.

While the original implementation of SAMtools and BCFtools provided APIs to parse the files, it mixed these APIs with application code. This did not guarantee long-term stability and made it difficult to interface in other programs. To solve this, in 2013 the decision was taken to separate the API from the command line tools and to produce HTSlib as a dedicated programming library that processes common data formats used in high-throughput sequencing. Support for the European Bioinformatics Institute's CRAM format was added and in 2014 the first official release (1.0) was published. This library implements stable and robust APIs that other programs can rely on. It enables efficient access to SAM/BAM/CRAM, VCF/BCF, FASTA, FASTQ, block-gzip compressed data, and indexes. It can be used natively in C/C++ code and has bindings to many other popular programming languages, such as Python, Rust, and R, boosting the development of sequence analysis tools.

HTSlib is not merely a separation; it also brought numerous improvements to SAMtools, BCFtools, and other third-party programs depending on it. HTSlib is linked into ~900 GitHub projects (see Supplementary Section S2), and HTSlib itself has been forked >300 times. HTSlib has been installed via bioconda >1 million times, and there are ~10,000 GitHub projects using it via Pysam. The library is freely available for commercial and non-commercial use (the MIT/BSD compatible license) from htslib.org and GitHub [13].

Findings

Implementation

The main purpose of HTSlib is to provide access to genomic information files, both alignment data (SAM, BAM, and CRAM formats) and variant data (VCF and BCF formats). The library also provides interfaces to access and index genome reference data in FASTA format and tab-delimited files with genomic coordinates.

Given the typical file sizes of genomic data, compression is necessary for efficient storage of data. HTSlib supports the GZIP-compatible format BGZF (Blocked GNU Zip Format), which limits the size of compressed blocks, thus allowing indexing and random access to the compressed files. HTSlib includes 2 stand-alone programs that work with BGZF; bgzip is a general-purpose compression tool while tabix works on tab-delimited genome coordinate files (e.g., BED and GFF) and provides indexing and random access. BGZF compression is also used for BAM, BCF,

and compressed FASTA files. The CRAM format uses column-specific compression methods including gzip, rANS [14], LZMA, and bzip2. The CRAM implementation in HTSlib learns the best-performing compression method on the fly (see Supplementary Section S3).

The HTSlib library is structured as follows: the media access layer (Fig. 1a) is a collection of low-level system and library (libcurl, knet) functions, which facilitate access to files on different storage environments (disk, memory, network) and over multiple protocols to various online storage providers (AWS S3, Google Cloud, GA4GH htsget [15]; Fig. 1b). This functionality is transparently available through a unified low-level stream interface hFILE (Fig. 1c). All file formats (Fig. 1d and e) are accessible through a higher-level file format-agnostic htsFILE interface, which provides functions to detect file types and set write options and provides common code for file iterators. Building on this layer are specializations for alignment (SAM, BAM, and CRAM) and variant (VCF and BCF) files and various auxiliary functions (Fig. 1f and g).

This API (Fig. 1f and g) can be roughly divided into several classes:

1. The File Access API has basic methods for opening and closing alignment and variant files, as well as reading and writing records to a file. HTSlib automatically determines the input file type by its contents and output type by filename. Further explicit control is provided for format, data layout (in CRAM), and file compression levels. Data sorted in genomic order may also be indexed at the time of writing (for alignment data) or at a later post-writing stage.
2. The Header API is a collection of methods that enables extensive control of SAM and VCF headers, including reading, writing, and parsing the header, accessing and updating individual tags, and adding and removing header lines.
3. The Data API provides methods for parsing, updating, and retrieving information from individual record fields on both alignment and variant data. The library also includes the ability to read multiple VCF and BCF files in parallel, transparently merging their contents, so that the reader can easily process records with matching genomic positions and alleles.
4. The Index/Iterator API offers the ability to extract information from the various index formats specific to genomic data (BAI, CSI, CRAI, TBI), and to create iterators for genomic files. The original BAI and TBI indices were limited to 512 Mb and were replaced by CSI allowing $\leq 2^{44}$ bp. Both sequence alignment and variant call formats have millions of records that can be indexed by genomic location. An iterator groups a list of target genomic regions into a list of file offsets and contains the stepping and filtering logic to allow the file reader to extract only the information of interest. Additionally, the library provides the `regidx` API, which allows the user to efficiently search and intersect regions from arbitrary row-oriented text formats.
5. The Mpileup API performs a data pivot. Alignment data in SAM, BAM, and CRAM are retrieved in row-oriented format, record by record. Data rotation (merging ≥ 1 input files) presents the sequence data in a column-oriented form per reference position. This information can be used for single-nucleotide polymorphism and indel calling, consensus generation, and to make alignment viewers. Mpileup can also optionally calculate base alignment quality scores for each read [16]. The base alignment quality scores can be used to reduce false-positive single-nucleotide polymorphism calls

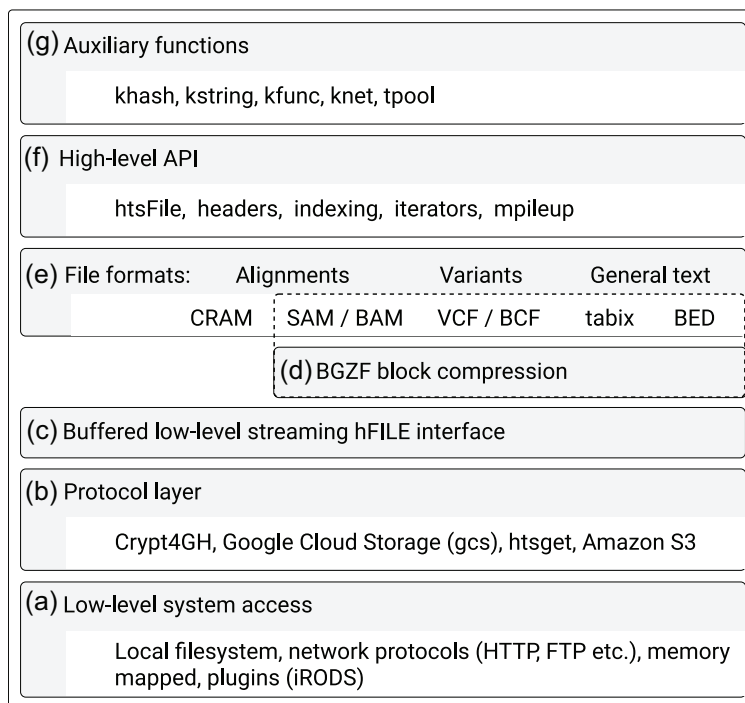


Figure 1: Htslib design.

by lowering the confidence scores at locations where the read alignment may be incorrect.

- HTSlib also includes various utility convenience functions such as hash tables, string manipulation, linked lists, heaps, sorting, and logging, and ensures portability between big- and little-endian platforms. Many of these originate from Klib [17]. A thread pool interface is provided for general multi-threading.

Benchmarks of sequence alignment formats

Given the widespread use of the library, performance and low memory requirements are paramount, which means that even relatively small improvements can lead to time and energy savings when analyzing large amounts of data.

To test maximum throughput for alignment data, elapsed times were obtained for each file type using both 1 main thread and also 16 additional worker threads. Not all tools supported indexing of all formats, and only in more recent HTSlib versions is there support for indexing and random access of BGZF-compressed SAM files. Full benchmarks are in Supplementary Table S5, with a summary for BAM shown here in Fig. 2.

The tests were performed on a RAM disk (/dev/shm) so represent maximum I/O rates for this system.

Figure 2 shows read and write elapsed timings for the BAM format on chromosome 1 of ENA accession ERR3239276. Note that “samtools 2013” refers to SAMtools version 0.1.19 and not the current SAMtools release. Other tool versions are HTSlib 1.10.2-32-ga22a0af, HTSJDK 2.22.0, Sambamba 0.7.1, and Scramble 1.14.13. These use up to 16 threads, but the HTSJDK times are with only 1 additional thread per file. SAMtools 0.1.19 has multi-threaded writing only, so the speed is limited by the reading portion. (See Supplementary Section S4 for full single-threaded and multi-threaded timings, along with benchmarks for the SAM and CRAM formats.)

HTSlib was the only tool tested capable of multi-threaded SAM decoding and encoding, which is important when processing output from a fast multi-threaded aligner. The use of the faster compression library libdeflate [18] over Zlib [19] is also a major contributing factor in BAM performance, meaning that BAM to BAM transcoding with 16 threads is 5 times faster than the original SAMtools 0.1.19 and BAM to SAM is 13 times faster.

File sizes also differ slightly for BAM, owing to differing Deflate implementations (Zlib, Libdeflate, and Intel deflate). HTSlib’s CRAM size is 24% smaller than HTSJDK, while being 4 times faster (with a single thread), although the files remain compatible (see Supplementary Table S5).

To compare the random access capabilities of HTSlib we chose gene and exon regions from the Ensembl database across chromosome 1 and measured the time and I/O statistics to retrieve all alignments overlapping those regions. HTSlib, HTSJDK, and Sambamba all support a multi-region iterator that is able to optimize I/O for many regions, reporting alignments that overlap multiple regions once only. SAMtools 0.1.19 and Scramble have no such feature; hence, regions that overlap will report some records multiple times and the same block may be read more than once.

The exon list had 58,160 regions (many overlapping each other) covering 5.5% of the chromosome. Figure 3 shows the random access efficiency, in both time and number of bytes read, for the exon list with BAM input. HTSlib is the fastest and requires less I/O to retrieve the same records.

Benchmarks of variant formats

The only common format supported between current HTSlib/BCFtools and HTSJDK is BGZF-compressed VCF. Figure 4 shows the time to read and read/write this format on a 929-sample test set [20] (see Supplementary Section S8). The source file is hgdp_wgs.20_190_516.full.chr20.vcf.gz, aligned and called

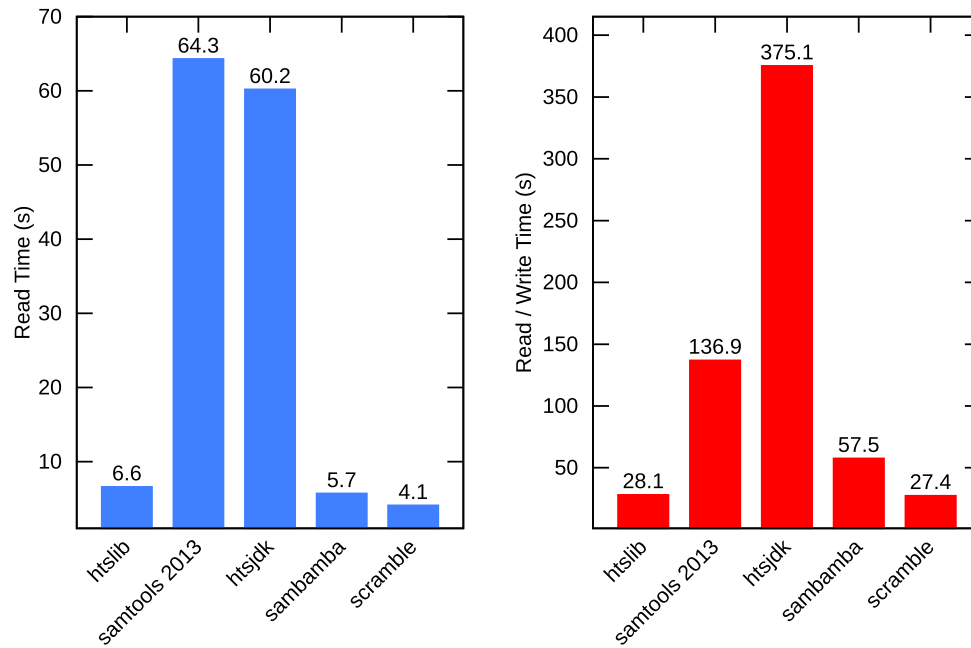


Figure 2: BAM elapsed read/write times, with up to 16 threads available.

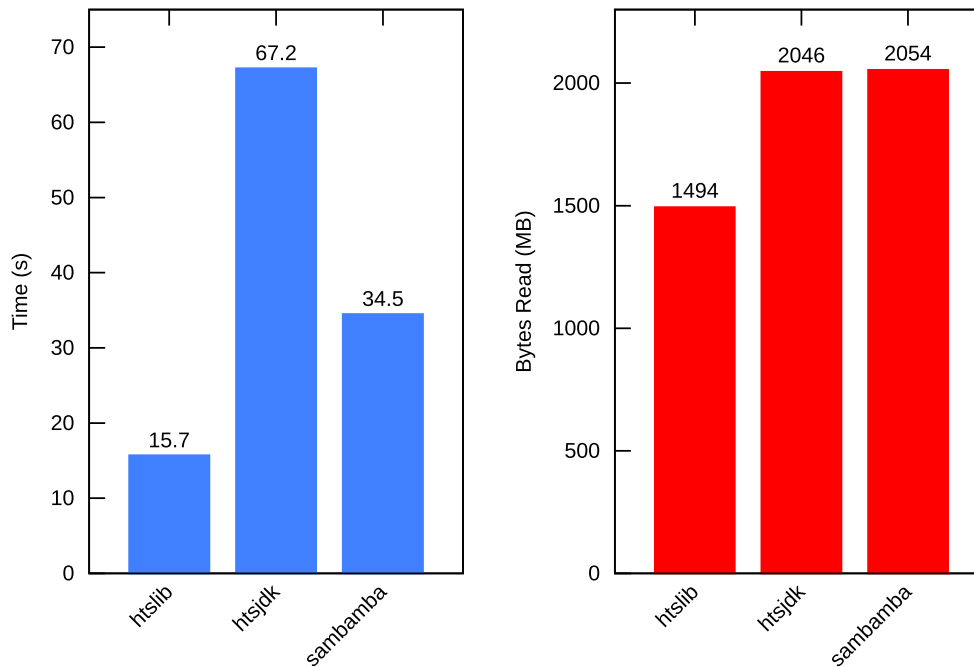


Figure 3: Time and megabytes of data read, for exon list in BAM.

from ENA PRJEB6463. Only single thread times are shown because currently multi-threading is suboptimal in BCFtools and not available in HTSJDK.

HTSlib also supports the BGZF-compressed BCF format, a binary variant of VCF. This performs considerably better than the compressed VCF, being 5 times faster to decode and nearly 3 times faster to encode (see Supplementary Tables S10 and S11 for details and more complete results).

Discussion

Over the lifetime of HTSlib the cost of sequencing has decreased by ~ 100 -fold, with a corresponding increase in data volume [21]. New sequencing technologies have also been developed that produce much longer reads. The alignment and variant file formats have moved on from being group-led research to being maintained by the File Formats subgroup of the Global Alliance

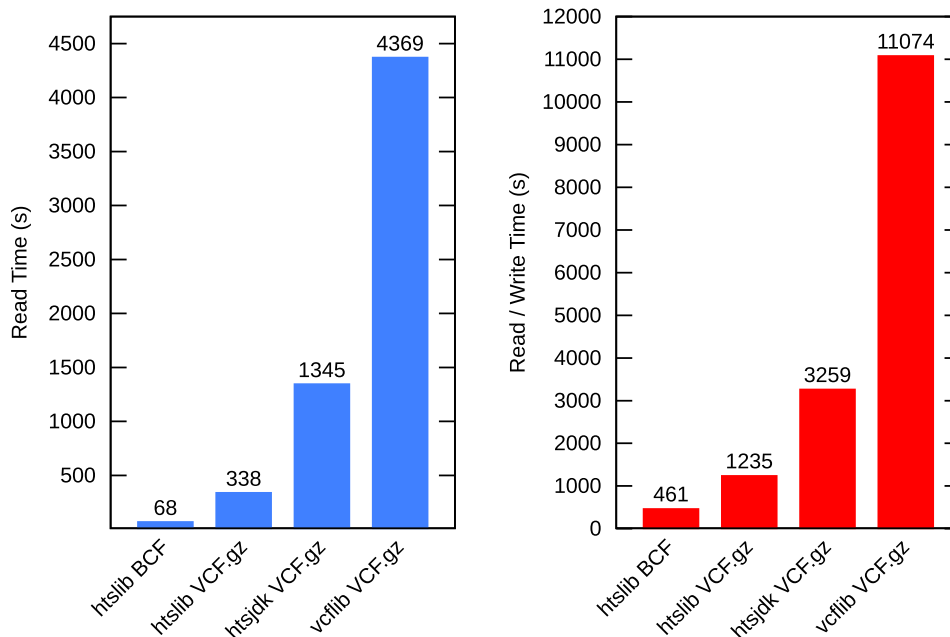


Figure 4: VCF.gz and BCF read and read/write times, 929 samples.

for Genomics and Health [22]. Over the years various improvements and modifications have been made to the specifications. Together these have been and will continue to be a driving force for continued development.

Since HTSlib 1.0, there have been 11 major releases and >1,600 code commits, more than doubling the number of lines of C code [15]. It has gained support for the CRAM file format, better indexing, extended APIs, more transfer protocols (S3, Google Cloud, Htsget), and improved threading and speed. Through the use of automated tests, static analysis tools, and fuzz testing it has been made much more reliable (see Supplementary Section S12).

Some of the existing limitations in HTSlib come from the design of the underlying file formats, e.g., BAM, CRAM, and BCF limit the maximum reference length to 2 Gb (see Supplementary Section S13). We expect future standards development to include improvements leading to better scaling of many-sample VCF, additional support for structural variation, better handling of very long sequencing reads, large genomes, and support for base modifications. Further plans include speeding up both the VCF parser and mpileup, improved documentation, and better support for BED files.

Availability of Supporting Source Code and Requirements

Project name: HTSlib

Project home page: <https://www.htslib.org>, <https://github.com/samtools/htslib>

Operating system(s): Platform independent

Programming language: C

License: A mix of Modified 2-Clause BSD (CRAM) and MIT/Expat (everything else)

RRID:SCR_002105

biotoools: htslib

Data Availability

The data set supporting the benchmarking results of this article is available in the ENA (ERR3239276 and PRJEB6463) and via FTP [23]. Snapshots of the code are also available via the GigaScience GigaDB repository [24].

Editors Note

An accompanying article describing the latest versions of SAM-tools and BCFtools is published alongside this article [25].

Additional Files

Supplementary Figure S1: Binary BCF vs VCF format

Supplementary Section S2: Estimated number of HTSlib source code clones

Supplementary Section S3: CRAM compression algorithm

Supplementary Section S4: Performance of HTSlib's SAM, BAM, CRAM implementations

Supplementary Table S5: Read and read/write timings for tools and file formats

Supplementary Table S6: Random access times and data volumes, single thread

Supplementary Table S7: Random access times and data volumes, 8 threads

Supplementary Section S8: Performance of HTSlib's VCF, BCF implementations

Supplementary Figure S9: VCF and BCF read and read/write speeds

Supplementary Table S10: Multi-sample VCF and BCF performance

Supplementary Table S11: Single-sample VCF and BCF performance

Supplementary Section S12: Automatic testing

Supplementary Section S13: The format size limitations

Abbreviations

API: Application Programming Interface; BAM: Binary sequence Alignment/Map; bp: base pairs; BCF: Binary variant Call Format; BGZF: Blocked GNU Zip Format; ENA: European Nucleotide Archive; Gb: gigabase pairs; Mb: megabase pairs; RAM: random access memory; SAM: Sequence Alignment/Map; VCF: Variant Call Format.

Competing Interests

The authors declare they have no competing interests.

Funding

This work was supported by the Wellcome Trust grant (206194).

Authors' contributions

J.K.B., P.D., R.M.D., H.L., J.M., V.O., and A.W. wrote the software. R.M.D., T.K., and J.M. supervised the project. J.K.B., P.D., R.M.D., and A.W. wrote the original draft of the manuscript with all authors reviewing.

References

1. The 1000 Genomes Project Consortium. A global reference for human genetic variation. *Nature* 2015;**526**: 68–74.
2. Li H, Handsaker B, Wysoker A, et al. The Sequence Alignment/Map format and SAMtools. *Bioinformatics* 2009;**25**:2078–9.
3. Danecek P, Auton A, Abecasis G, et al. The variant call format and VCFtools. *Bioinformatics* 2011;**27**:2156–8.
4. Li H. A statistical framework for SNP calling, mutation discovery, association mapping and population genetical parameter estimation from sequencing data. *Bioinformatics* 2011;**27**:2987–93.
5. SAMtools, Version 1.11 2020. <https://github.com/samtools/samtools/releases/tag/1.11>. Accessed 9 February 2021.
6. BCFtools, Version 1.11. 2020. <https://github.com/samtools/bcftools/releases/tag/1.11>. Accessed 9 February 2021.
7. HTSJDK, Version 2.23.0. 2020. <https://github.com/samtools/htsjdk>. Accessed 4 May 2020.
8. Tarasov A, Vilella AJ, Cuppen E, et al. Sambamba: fast processing of NGS alignment formats. *Bioinformatics* 2015;**31**:2032–4.
9. Bonfield JK. The Scramble conversion tool. *Bioinformatics* 2014;**30**:2818–9.
10. Hsi-Yang Fritz M, Leinonen R, Cochrane G, et al. Efficient storage of high throughput DNA sequencing data using reference-based compression. *Genome Res* 2011;**21**:734–40.
11. CRAM homepage. 2020. <https://www.ga4gh.org/cram/>. Accessed 6 November 2020.
12. Buels R, Yao E, Diesh CM, et al. JBrowse: a dynamic web platform for genome visualization and analysis. *Genome Biol* 2016;**17**:66, doi:10.1186/s13059-016-0924-1.
13. SAMtools. HTSlib, Version 1.11. 2020. <https://github.com/samtools/htslib/releases/tag/1.11>. Accessed 9 February 2021.
14. Duda J. Asymmetric numeral systems: entropy coding combining speed of Huffman coding with compression rate of arithmetic coding. *arXiv* 2013: 1311.2540.
15. Kelleher J, Lin M, Albach CH, et al. htsget: a protocol for securely streaming genomic data. *Bioinformatics* 2019;**35**:119–21.
16. Li H. Improving SNP discovery by base alignment quality. *Bioinformatics* 2011;**27**:1157–8.
17. Li H. Klib. 2013. <https://github.com/attractivechaos/klib>. Accessed 27 November 2020.
18. Biggers E. libdeflate, Version 1.7. 2020. <https://github.com/ebiggers/libdeflate>. Accessed 6 November 2020.
19. Deutsch P, Gailly J-L. ZLIB Compressed Data Format Specification version 3.3. 1996. <https://doi.org/10.17487/RFC1950>.
20. Bergström A, McCarthy SA, Hui R, et al. Insights into human genetic variation and population history from 929 diverse genomes. *Science* 2020;**367**:eaay5012.
21. National Human Genome Research Institute. DNA Sequencing Costs: Data. 2020. <https://www.genome.gov/about-genomics/fact-sheets/DNA-Sequencing-Costs-Data>. Accessed 23 September 2020.
22. Birney E, Vamathevan J, Goodhand P. Genomics in healthcare: GA4GH looks to 2022. *bioRxiv* 2017: doi.org/10.1101/203554.
23. Source file for hgdp.wgs.20190516.full.chr20. ftp://ngs.sanger.ac.uk/production/hgdp/hgdp_wgs.20190516/hgdp_wgs.20190516.full.chr20.vcf.gz. Accessed 27 November 2020.
24. Bonfield JK, Marshall J, Danecek P, et al. Supporting data for “HTSlib - C library for reading/writing high-throughput sequencing data.” *GigaScience Database* 2021. <http://dx.doi.org/10.5524/100867>.
25. Danecek P, Bonfield JK, Liddle J, et al. Twelve years of SAMtools and BCFtools. *GigaScience* 2021;**10**(2): doi:10.1093/gigascience/giab008