

The 1000 Genomes Project: data management and community access

Laura Clarke¹, Xiangqun Zheng-Bradley¹, Richard Smith¹, Eugene Kulesha¹, Chunlin Xiao², Iliana Toneva¹, Brendan Vaughan¹, Don Preuss², Rasko Leinonen¹, Martin Shumway², Stephen Sherry², Paul Flicek¹ & The 1000 Genomes Project Consortium³

The 1000 Genomes Project was launched as one of the largest distributed data collection and analysis projects ever undertaken in biology. In addition to the primary scientific goals of creating both a deep catalog of human genetic variation and extensive methods to accurately discover and characterize variation using new sequencing technologies, the project makes all of its data publicly available. Members of the project data coordination center have developed and deployed several tools to enable widespread data access.

High-throughput sequencing technologies, including those from Illumina, Roche Diagnostics (454) and Life Technologies (SOLiD), enable whole-genome sequencing at an unprecedented scale and at dramatically reduced costs over the gel capillary technology used in the human genome project. These technologies were at the heart of the decision in 2007 to launch the 1000 Genomes Project, an effort to comprehensively characterize human variation in multiple populations. In the pilot phase of the project, the data helped create an extensive population-scale view of human genetic variation¹.

The larger data volumes and shorter read lengths of high-throughput sequencing technologies created substantial new requirements for bioinformatics, analysis and data-distribution methods. The initial plan for the 1000 Genomes Project was to collect 2× whole genome coverage for 1,000 individuals, representing ~6 giga-base pairs of sequence per individual and ~6 tera-base pairs (Tbp) of sequence in total. Increasing sequencing capacity led to repeated revisions of these plans to the current project scale of collecting low-coverage, ~4× whole-genome and ~20× whole-exome sequence for ~2,500 individuals plus high-coverage,

~40× whole-genome sequence for 500 individuals in total (~25-fold increase in sequence generation over original estimates). In fact, the 1000 Genomes Pilot Project collected 5 Tbp of sequence data, resulting in 38,000 files and over 12 terabytes of data being available to the community¹. In March 2012 the still-growing project resources include more than 260 terabytes of data in more than 250,000 publicly accessible files.

As in previous efforts^{2–4}, the 1000 Genomes Project members recognized that data coordination would be critical to move forward productively and to ensure the data were available to the community in a reasonable time frame. Therefore, the Data Coordination Center (DCC) was set up jointly between the European Bioinformatics Institute (EBI) and the National Center for Biotechnology (NCBI) to manage project-specific data flow, to ensure archival sequence data deposition and to manage community access through the FTP site and genome browser.

Here we describe the methods used by the 1000 Genomes Project members to provide data resources to the community from raw sequence data to project results that can be browsed. We provide examples drawn from the project's data-processing methods to demonstrate the key components of complex workflows.

Data flow

Managing data flow in the 1000 Genomes Project such that the data are available within the project and to the wider community is the fundamental bioinformatics challenge for the DCC (**Fig. 1** and **Supplementary Table 1**). With nine different sequencing centers and more than two dozen major analysis groups¹, the most important initial challenges are (i) collating all

¹European Bioinformatics Institute, Wellcome Trust Genome Campus, Hinxton, Cambridge, UK. ²National Center for Biotechnology Information, National Library of Medicine, US National Institutes of Health, Bethesda, Maryland, USA. ³A full list of members is provided in the **Supplementary Note**. Correspondence should be addressed to P.F. (flicek@ebi.ac.uk) or The 1000 Genomes Project Consortium (info@1000genomes.org).

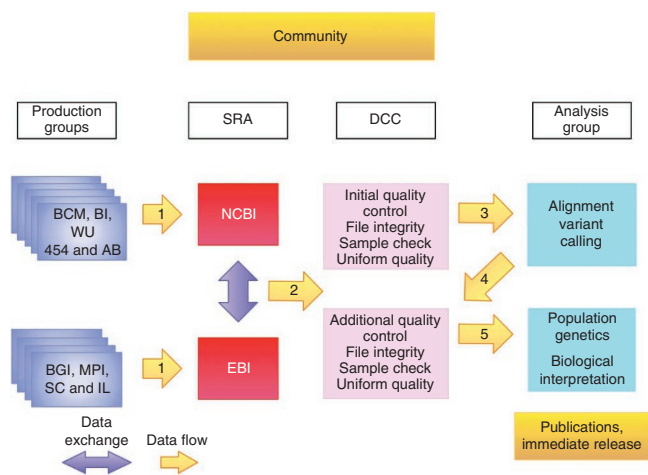


Figure 1 | Data flow in the 1000 Genomes Project. The sequencing centers submit their raw data to one of the two SRA databases (arrow 1), which exchange data. The DCC retrieves FASTQ files from the SRA (arrow 2) and performs quality-control steps on the data. The analysis group access data from the DCC (arrow 3), aligns the sequence data to the genome and uses the alignments to call variants. Both the alignment files and variant files are provided back to the DCC (arrow 4). All the data are publicly released as soon as possible. BCM, Baylor College of Medicine; BI, Broad Institute; WU, Washington University; 454, Roche; AB, Life Technologies; MPI, Max Planck Institute for Molecular Genetics; SC, Wellcome Trust Sanger Institute; IL, Illumina.

the sequencing data centrally for necessary quality control and standardization; (ii) exchanging the data between participating institutions; (iii) ensuring rapid availability of both sequencing data and intermediate analysis results to the analysis groups; (iv) maintaining easy access to sequence, alignment and variant files and their associated metadata; and (v) providing these resources to the community.

In recent years, data transfer speeds using TCP/IP-based protocols such as FTP have not scaled with increased sequence production capacity. In response, some groups have resorted to sending physical hard drives with sequence data⁵, although handling data this way is very labor-intensive. At the same time, data-transfer requirements for sequence data remain well below those encountered in physics and astronomy, so building a dedicated network infrastructure was not justified. Instead, the project members elected to rely on an internet transfer solution from the company Aspera, a UDP-based method that achieves data transfer rates 20–30 times faster than FTP in typical usage. Using Aspera, the combined submission capacity of the EBI and NCBI currently approaches 30 terabytes per day, with both sites poised to grow as global sequencing capacity increases.

The 1000 Genomes Project was responsible for the first multi-terabase submissions to the two sequence-read archives (SRAs): the SRA at the EBI, provided as a service of the European Nucleotide Archive (ENA), and the NCBI SRA⁶. Over the course of the project, the major sequencing centers developed automated data-submission methods to either the EBI or the NCBI, whereas both SRA databases developed generalized

methods to search and access the archived data. The data formats accepted and distributed by both the archives and the project have also evolved from the expansive sequence read format (SRF) files to the more compact Binary Alignment/Map (BAM)⁷ and FASTQ formats (Table 1). This format shift was made possible by a better understanding of the needs of the project-analysis group, leading to a decision to stop archiving raw intensity measurements from read data to focus exclusively on base calls and quality scores.

As a ‘community resource project’⁸, the 1000 Genomes Project publicly releases prepublication data as described below as quickly as possible. The project has mirrored download sites at the EBI (<ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/>) and NCBI (<ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/>) that provide project and community access simultaneously and efficiently increase the overall download capacity. The master copy is directly updated by the DCC at the EBI, and the NCBI copy is usually mirrored within 24 h via a nightly automatic Aspera process. Generally users in the Americas will access data most quickly from the NCBI mirror, whereas users in Europe and elsewhere in the world will download faster from the EBI master.

The raw sequence data, as FASTQ files, appear on the 1000 Genomes FTP site within 48–72 h after the EBI SRA has processed them. This processing requires that data originally submitted to the NCBI SRA must first be mirrored at the EBI. Project data are managed through periodic data freezes associated with a dated `sequence.index` file (Supplementary Note). These files were produced approximately every two months during the pilot phase, and for the full project the release frequency varies depending on the output of the production centers and the requirements of the analysis group.

Alignments based on a specific `sequence.index` file are produced within the project and distributed via the FTP site in BAM format, and analysis results are distributed in variant call format (VCF) format⁹. Index files created by the Tabix software¹⁰ are also provided for both BAM and VCF files.

All data on the FTP site have been through an extensive quality-control process. For sequence data, this includes checking syntax and quality of the raw sequence data and confirming sample identity. For alignment data, quality control includes file integrity and metadata consistency checking (Supplementary Note).


Data access

The entire 1000 Genomes Project data set is available, and the most logical approach to obtain it is to mirror the contents of the FTP site, which is, as of March 2012, more than 260 terabytes.

Table 1 | File formats used in the 1000 Genomes Project

Format	Description	Website
SRF	Container format for data from sequencing machines based on ZTR	http://srf.sourceforge.net/
FASTQ	Text-based format for sequence and quality values	http://en.wikipedia.org/wiki/FASTQ_format
SAM or BAM ⁷	A compact alignment format for placement of short-read data with respect to a reference genome (designed by consortium members)	http://samtools.sourceforge.net/
VCF ⁹	A column-based text format for storing variant calls and individual genotypes for those calls (designed by consortium members)	http://vcftools.sourceforge.net/

Figure 2 | Remote file viewing. The 1000 Genomes Browser enables the attachment of remote files to allow accessible BAM and VCF files to be displayed in 'Location' view. The tracks in the image from our October 2011 browser based on Ensembl version 63 are an NA12878 BAM file from EBI FTP site with consensus sequence noted by the upper arrow and sequence reads by the lower arrow (i); variants from 20110521 release VCF file shown as a track with two variants in yellow (ii); variants from the 20101123 release database shown as a track with one variant in yellow (iii); and gene annotation from Ensembl showing the genomic context (iv). The ability for users to view data from files allows rapid access to new data before the database can be updated.



Our experience is that most users are more interested in analysis results and targeted raw data or alignment slices from specific regions of the genome rather than the entire data set. Indeed, the analysis files are distributed via the FTP site in directories named for the sequence.index freeze date they are based on (**Supplementary Note**). However, with hundreds of thousands of available files, locating and accessing specific project data by browsing the FTP directory structure can be extremely difficult.

A file called *current.tree* is provided at the root of the FTP site to assist in searching the site. This file was designed to enable mirroring the FTP site and contains a complete list of all files and directories including time of last update and file integrity information. We developed a web interface (<http://www.1000genomes.org/ftpsearch/>) to provide direct access to the *current.tree* file using any user-specified sample identifier(s) or other information found in our data file names, which follow a strict convention to aid searching. The search returns full file paths to either the EBI or the NCBI FTP site and supports filters to exclude file types likely to produce a large number of results such as FASTQ or BAM files (**Supplementary Note**).

For users wanting discovered variants or alignments from specific genomic regions without downloading the complete files, they can obtain subsections of BAM and VCF files either directly with Tabix or via a web-based data-slicing tool (**Supplementary Note**). VCF files can also be divided by sample name or population using the data slicer.

One can view 1000 Genomes data in the context of extensive genome annotation, such as protein-coding genes and whole-genome regulatory information through the dedicated 1000 Genomes browser based on the Ensembl infrastructure¹¹ (<http://browser.1000genomes.org/>). The browser displays project variants before they are processed by dbSNP or appear in genome resources such as Ensembl or the University of California Santa Cruz (UCSC) genome browser. The 1000 Genomes browser also provides Ensembl variation tools including the Variant Effect Predictor (VEP)¹² as well as 'sorting tolerant from intolerant' (SIFT)¹³ and PolyPhen¹⁴ predictions for all nonsynonymous variants (**Supplementary Note**). The browser supports viewing of both 1000 Genomes Project and other web-accessible indexed BAM and VCF files in genomic context (**Fig. 2**). A stable archival

version of the 1000 Genomes browser based on Ensembl code release 60 and containing the pilot project data is available at <http://pilotbrowser.1000genomes.org/>.

The underlying MySQL databases that support the project browser are also publicly available and these can be directly queried or accessed programmatically using the appropriate version of the Ensembl Application Programming Interface (API) (**Supplementary Note**).

Users may also explore and download project data using the NCBI data browser at <http://www.ncbi.nlm.nih.gov/variation/tools/1000genomes/>. The browser displays both sequence reads and individual genotypes for any region of the genome. Sequence for selected individuals covering the displayed region can be downloaded in BAM, SAM, FASTQ or FASTA format. Genotypes can likewise be downloaded in VCF format (**Supplementary Note**).

The project submits all called variants to the appropriate repositories using the handle "1000GENOMES". Pilot project single-nucleotide polymorphisms and small indels have been submitted to dbSNP¹⁵, and structural variation data have been submitted to the Database of Genomic Variants archive (DGVA)¹⁶. Full project variants will be similarly submitted.

For users of Amazon Web Services, all currently available project BAM and VCF files are available as a public data set via <http://1000genomes.s3.amazonaws.com/> (**Supplementary Note**).

Finally, all links and announcements of the project data can be found at <http://www.1000genomes.org/>, and announcements are made available via rss (<http://www.1000genomes.org/announcements/rss.xml>), Twitter (@1000genomes) and also via an email list (1000announce@1000genomes.org; **Supplementary Note**).

Discussion

Methods of data submission and access developed to support the 1000 Genomes Project offer benefits to all large-scale sequencing projects and the wider community. The streamlined archival process takes advantage of the two synced copies of the SRA, which

distribute the resource-intensive task of submission processing. In addition, the close proximity of the DCC to the SRA ensures that all 1000 Genomes data are made available to the community as quickly as possible and allowed the archives to benefit from the lessons learned by the DCC.

Large-scale data generation and analysis projects can benefit from an organized and centralized data-management activity²⁻⁴. The goals of such activities are to provide necessary support and infrastructure to the project while ensuring that data are made available as rapidly and widely as possible. In supporting the 1000 Genome Project analysis, the established extensive data flow includes multiple tests to ensure data integrity and quality (Fig. 1). As part of this process, data are made available to members of the consortium and the public simultaneously at specific points in the data flow, including at the collection of sequence data and the completion of alignments.

Beyond directly supporting the needs of the project, centralized data management ensures that resources targeted to users outside the consortium analysis group are created. These include the 1000 Genomes Browser (<http://browser.1000genomes.org/>), submission of both preliminary and final variant data sets to dbSNP and to dbVar/DGVA, provisioning of alignment and variant files in the Amazon Web Services cloud, and centralized variation annotation services.

The experiences of data management used for this project reflect in part the difficulty of adopting existing bioinformatics systems to new technologies and in part the challenge of data volumes much larger than previously encountered. The rapid evolution of analysis and processing methods is indicative of the community effort to provide effective tools for understanding the data.

Note: Supplementary information is available on the Nature Methods website.

ACKNOWLEDGMENTS

For early work and support to the DCC, we thank Z. Iqbal, H. Khouri, F. Cunningham, Y. Chen, W. McLaren, V. Zalunin, R. Radhakrishnan, D. Smirnov, J. Paschall, Z. Belaia, R. Sanders, C. O'Sullivan, S. Keenan, G. Ritchie and G. Cochrane. For maintenance of the EBI computer infrastructure, we acknowledge J. Barker, V. Silventoinen, G. Kellman and P. Jokinen. Funding support at the EBI is provided by the Wellcome Trust (grant WT085532) and the European Molecular Biology Laboratory. This research was supported in part by the Intramural Research Program of the US National Institutes of Health National Library of Medicine.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Published online at <http://www.nature.com/naturemethods/>.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported (CC BY-NC-SA) license. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/3.0/>.

- 1000 Genomes Project Consortium. A map of human genome variation from population-scale sequencing. *Nature* **467**, 1061–1073 (2010).
- Thorisson, G.A., Smith, A.V., Krishnan, L. & Stein, L.D. The International HapMap Project Web site. *Genome Res.* **15**, 1592–1593 (2005).
- Rosenbloom, K.R. *et al.* ENCODE whole-genome data in the UCSC Genome Browser. *Nucleic Acids Res.* **38**, D620–D625 (2010).
- Washington, N.L. *et al.* The modENCODE Data Coordination Center: lessons in harvesting comprehensive experimental details. *Database (Oxford)* **2011**, bar023 (2011).
- Baker, M. Next-generation sequencing: adjusting to data overload. *Nat. Methods* **7**, 495–499 (2010).
- Shumway, M., Cochrane, G. & Sugawara, H. Archiving next generation sequencing data. *Nucleic Acids Res.* **38**, D870–D871 (2010).
- Li, H. *et al.* The Sequence Alignment/Map format and SAMtools. *Bioinformatics* **25**, 2078–2079 (2009).
- Toronto International Data Release Workshop Authors. Prepublication data sharing. *Nature* **461**, 168–170 (2009).
- The Toronto Agreement describes a set of best practices for prepublication data sharing. These practices have been adopted by the 1000 Genomes Project and have helped drive the widespread use of the data.**
- Danecek, P. *et al.* The variant call format and VCFtools. *Bioinformatics* **27**, 2156–2158 (2011).
- Li, H. Tabix: fast retrieval of sequence features from generic TAB-delimited files. *Bioinformatics* **27**, 718–719 (2011).
- Flicek, P. *et al.* Ensembl 2011. *Nucleic Acids Res.* **39**, D800–D806 (2011).
- McLaren, W. *et al.* Deriving the consequences of genomic variants with the Ensembl API and SNP Effect Predictor. *Bioinformatics* **26**, 2069–2070 (2010).
- The Ensembl Variant Effect Predictor (VEP) is a flexible and regularly updated method to annotate all newly discovered variants and provides information about how such variants impact genes, regulatory regions and other key genomic features.**
- Kumar, P., Henikoff, S. & Ng, P.C. Predicting the effects of coding non-synonymous variants on protein function using the SIFT algorithm. *Nat. Protoc.* **4**, 1073–1081 (2009).
- Adzhubei, I.A. *et al.* A method and server for predicting damaging missense mutations. *Nat. Methods* **7**, 248–249 (2010).
- Foelø, M.L. & Sherry, S.T. NCBI dbSNP database: content and searching. in *Genetic Variation: A Laboratory Manual* (eds., Weiner, M.P., Gabriel, S.B. & Stephens, J.C.) 41–61 (Cold Spring Harbor Laboratory Press, 2007).
- Church, D.M. *et al.* Public data archives for genomic structural variation. *Nat. Genet.* **42**, 813–814 (2010).