
Benchmarking Database Systems for Genomic Selection Implementation

Yaw Nti-Addae^{1*}, Dave Matthews², Victor Jun Ulat³, Raza Syed¹, Guilhem Sempéré⁴, Adrien Pétel⁵, Jon Renner⁶, Pierre Larmande⁷, Valentin Guignon⁸, Elizabeth Jones¹, Kelly Robbins⁹

¹Department of Biotechnology, Cornell University, ²Boyce Thompson Institute, ³Centro Internacional de Mejoramiento de Maíz y Trigo (CIMMYT), ⁴INTERTRYP, Univ Montpellier, CIRAD, IRD, ⁵UMR PVBMT, CIRAD, ⁶University of Minnesota, ⁷UMR DIADE, IRD, University of Montpellier, ⁸Bioversity International, ⁹College of Agriculture and Life Sciences, Cornell University

*To whom correspondence should be addressed.

Abstract

Motivation: With high-throughput genotyping systems now available, it has become feasible to fully integrate genotyping information into breeding programs [22]. To make use of this information effectively requires DNA extraction facilities and marker production facilities that can efficiently deploy the desired set of markers across samples with a rapid turnaround time that allows for selection before crosses needed to be made. In reality, breeders often have a short window of time to make decisions by the time they are able to collect all their phenotyping data and receive corresponding genotyping data. This presents a challenge to organize information and utilize them in downstream analyses to support decisions made by breeders. In order to implement genomic selection routinely as part of breeding programs one would need an efficient genotype data storage system. We selected and benchmarked six popular open-source data storage systems, including relational database management and columnar storage systems.

Results: We found that data extract times are greatly influenced by the orientation in which genotype data is stored in a system. HDF5 consistently performed best, in part because it can more efficiently work with both orientations of the allele matrix.

Availability: <http://gobiin1.bti.cornell.edu:6083/projects/GBM/repos/benchmarking/browse>

Contact: yn259@cornell.edu

1 Introduction

The development of Next-generation sequencing (NGS) technologies has made it feasible to generate huge volumes of genomic data. In the field of plant breeding, the availability of cost-effective genomic data has the potential to change the way crop breeding is done. Genomic selection (GS) is a breeding method where the performance of new plant varieties is predicted based on genomic information [1]. Multiple studies have shown the potential of this methodology to increase the rates of genetic gain in breeding programs by decreasing generation interval, the time it takes to screen new offspring and identify the best performers for use as parents in the next generation [2], [3]. Although model capabilities exist for the implementation of GS, mainstream applications require the computational infrastructure to manage NGS data, often generated from highly multiplexed sequencing runs that generate low coverage data with large amounts of missing information. The lack of computational infrastructure remains a major barrier to routine use of genomic information in public sector breeding programs.

The Genomic Open-source Breeding Informatics Initiative (GOBii) is a project aimed at increasing the rates of genetic gain in crop breeding programs serving regions in Africa and South Asia by developing the capabilities required for routine use of genomic information. To deal with the technical challenges of storage, rapid access (i.e. query execution), and computation on NGS data, the initial focus of GOBii has been the development of an efficient genomic data management system (GOBii-GDM). The system must be able to efficiently store huge volumes of genomic information and provide rapid data extraction for computation. The system must be scalable for large breeding programs, while being able to run effectively at institutions with limited access to large computational clusters. While many open-source technologies exist for the management of large two-dimensional datasets, it is unclear which technologies best suit the needs of plant breeding and genetics research.

There are many appealing characteristics of traditional relational database management systems (RDBMS), which are designed and built to store, manage, and analyze large-scale data. However, performance can be problematic when dealing with large matrix data like those commonly encountered in genomic research. One common limitation of RDBMS is database partitioning, which allows for a logical database to be divided into constituent parts and distributed over a number of nodes, (e.g. in a

computer cluster) [4]. To address these performance issues, many RDBMS have capabilities for working with binary large objects (BLOBs). Current versions of PostgreSQL (version 9.3 and up) have support for JSONB objects that could be used to store BLOBs of genomic data. However, this still does not solve the data retrieval performance issues [4]. An alternative to storing genomic data directly in a RDBMS is to use a hybrid system [5] with high-dimensional data being stored in files and key meta information required for querying the data stored in an RDBMS.

Leveraging several decades of work of the database community on optimizing query processing, columnar store databases, such as MonetDB are designed to provide high performance on complex queries against large databases, such as combining tables with hundreds of columns and millions of rows. In the biological context, experiences show that MonetDB [8] enables the data to be stored in a format to allow fast queries of vectors of genomic data based on marker or sample indexes, which should improve performance relative to RDBMS. More recently NoSQL systems have emerged as effective tools for managing high-dimensional genomic data [9], [10], [11]. NoSQL systems for distributed file storage and searching represent scalable solutions compare to RDBMS when dealing with semi-structured data types [12], [13], and MongoDB, a document-based NoSQL database has been used to develop a web-based tool for exploring genotypic information [14].

The Hierarchical Data File Format (HDF5) is a member of the high-performance distributed file systems family. It is designed for flexible, efficient I/O and for high-volume and complex data. It has demonstrated superior performance with high-dimensional and highly structured data such as genomic sequencing data [6] making it an appealing option for a hybrid system approach. There are an increasing number of bioinformatics applications, such as BioHDF [20], SnpSeek [7], Oxford Nanopore PoreTools [21] and FAST5, all of which use HDF5 to scale up simple queries execution on large number of documents. However, there is little reported information on the performance of HDF5 when the system is used to process more complex analytical queries that involve aggregations and joins.

To determine the ideal technology to serve as the backend of the GO-Bii-GDM, testing was performed using a large genotype-by-sequencing (GBS) dataset [15], [16], [19]. Open-source RDBMS, PostgreSQL and MariaDB, a community-developed fork under the GNU GPL of MySQL, were used as a baseline for performance testing and compared with HDF5, MonetDB, Elasticsearch [17], Spark [18], and MongoDB. Loading and extraction times were measured using queries that would be commonly run for applications of GS in a breeding program.

2 Methods

Six database systems were tested using a subset of a maize nested association mapping (NAM) population GBS SNP dataset [19] containing allele calls for 5258 samples (germplasm lines) and 31,617,212 markers. Each marker represents a physical position in the reference genome where polymorphisms were detected in the samples. Each genotyping call was encoded as a single ASCII character for the diploid result, using IUPAC nucleotide ambiguity codes for heterozygotes and “N” for missing data. The input and output format for all tests was a text file containing only the tab-delimited allele calls.

(a) (b)

Genomic matrices can be stored in two different orientations as showing in Figure 1. Given that traditional RDBMS (PostgreSQL and MariaDB) or disk stores are optimized for extracting data by rows, and

columnar stores (MonetDB and Parquet) are optimized for extracting data




Figure 1: Different orientation of genotyping data. (a) markers in rows and samples in columns, whereas (b) shows markers in columns and samples in rows

by columns, we stored data in these two different orientations where possible for each of the database systems. Unfortunately, due to size and technology restricts, not all systems could handle both orientations. We tested if the orientation of genotype matrix has an impact on query execution times. For the purposes of this benchmarking, we will define “marker-fast” orientation as the orientation of the genotype matrix in a system that favors the extract of markers, and converse, “sample-fast” as orientation of data in a system that favors the extract of samples. For example, in MonetDB, sample-fast orientation will have samples in columns and markers in rows and as such favor the querying of samples to markers. Vice versa, sample-fast orientation in PostgreSQL will have samples in rows and markers as indexes in a binary JSON object.

Three use cases were used to test the performance of systems with queries set up to extract data by:

- I. All samples for a list of markers (USECASE I)
- II. All markers for a list of samples (USECASE II)
- III. A block of data defined by a list of markers and samples (USECASE III)

For each use case, we tested extracting a contiguous list of markers or samples versus a random list. Care was necessary to avoid unrepeatability results due to memory caches in the system. All tests were run on a single node server with a 10 gigabit ethernet connection to a fast-access file server. The server specifications and version of the tested systems are listed in Table 1.

Table 1. Server configuration and software versions.

Server Configuration	
Processor	Intel Xeon E5 2620 V3 2.4GHz 6C 15Mb
Number of Processors	24
Memory	128GB DDR-4 2133MHz ECC/REG
Raid	6
Network File Storage	50TB
Operating System	Debian
Software	Version
HDF5	1.8
MariaDB	10.1
MonetDB	1.1
MongoDB	3.6.0 (compression level : snappy)

Benchmarking database systems for GS

PostgreSQL	9.5
Spark	2.3
Java Hotspot	1.8 (64 bit)
Elasticsearch	2.3.3 (under Java 1.8.0_92 64bit)

Database Implementation

The default parameters were applied to each benchmarked system, with some parameters critical for performance, such as memory allocation, manually optimized.

PostgreSQL Implementation

PostgreSQL is one of the most widely used open source object-relational database systems. Tests were done in version 9.5 and its configuration file was modified to consume up to 30GB of memory per query, whereas the default configurations use only 250MB of memory. Data was stored in “marker-fast” orientation for USECASE I as shown in Table 2.

Marker	Data*
S6_120018	{"B73 (PI550473):250028951": "C", "B97 (PI564682):250027958": "T", "CML103 (Ames27081):250027883": "C", "CML228 (Ames27088):250027960": "C"}
S6_120046	{"B73 (PI550473):250028951": "A", "B97 (PI564682):250027958": "A", "CML103 (Ames27081):250027883": "A", "CML228 (Ames27088):250027960": "A"}

*Data stored as JSONB objects

Conversely, data was stored in sample-fast orientation for USECASE II, where samples were in one column and associated markers and SNP allele calls in JSONB format in another column. We recognize that PostgreSQL can be optimized in cases of highly sparse data by storing only non “N” SNPs, which allows PostgreSQL to act like a document store, thereby reducing the size of stored information by many folds. For the purposes of this benchmarking, all data was treated as complete as many USECASES in genomic selection will pull information in which missing data has been imputed.

MariaDB Implementation

MariaDB is a community-developed, commercially supported fork of MySQL relational database. MariaDB was configured to utilize up to the maximum available memory on the server. Similar to PostgreSQL, data was stored in marker-fast for USECASE I, as shown in Table 3, and sample-fast for USECASE II. Dynamic columns allow one to store different sets of columns for each row in a table. It works by storing a set of columns in a BLOB and having a small set of functions to manipulate it.

Marker	Data*
S6_120018	{"B73 (PI550473):250028951": "C", "B97 (PI564682):250027958": "T", "CML103 (Ames27081):250027883": "C", "CML228 (Ames27088):250027960": "C"}
S6_120046	{"B73 (PI550473):250028951": "A", "B97 (PI564682):250027958": "A", "CML103 (Ames27081):250027883": "A", "CML228 (Ames27088):250027960": "A"}

*Data is stored as dynamic columns

MongoDB Implementation

MongoDB is an open source document-oriented database system. MongoDB tests were performed using version 3.6.0 configured with the WiredTiger storage engine and the snappy compression level, a choice driven by a comparison work done in [14]. MongoDB was tested with data stored in both orientations, marker-fast and sample-fast. For the marker-fast orientation, two types of documents were used:

- one for storing the genotype array corresponding to each marker, as follows:

```
{
  "_id": "S6_120018",
  "g": [
    "T",
    "A",
    "C",
    "N", ... ]
}
```

- the second for mapping sample names to indices in the latter array

```
{
  "_id": "CML247(PI595541):250027894",
  "n": NumberInt(0)
}
```

In the sample orientation, the sample collection remained the same as above, but the documents storing genotypes were refactored as:

```
{
  "_id": {
    "ii": "B73(PI550473):250028951",
    "c": NumberInt(0)
  },
  "g": [
    "C",
    "A",
    "G",
    "C", ... ]
}
```

Ideally, the document id would have been just the sample name, and the genotype array length would have been equal to the number of markers, i.e. 31,617,212. But because MongoDB has a 16Mb document-size limitation, we had to split each sample’s genotype array into chunks of maximum size 100,000. This explains why the id here is composite and consists in the following pair: sample name + chunk index.

HDF5 Implementation

Hierarchical Data Format (HDF5) file format is designed to store and organize extremely large and complex data collections. The allele matrix was stored in one-byte cells in both orientations, marker-fast with samples as columns (HDF5 dimension 1) and markers as rows (dimension 0) and the other in the opposite orientation. When extracting data for a list of samples and a list of markers, USECASE III, the most straightforward approach would be to use HDF5’s *H5Sselect_elements* function to address each allele call by its (sample, marker) coordinates. However, it was found to be much faster to extract all the data for each marker into a memory array using *H5Sselect_hyperslab*, and then look up the results for the desired samples in that array. The speed increase from this approach was more than 30-fold under all conditions tested.

Under some conditions HDF5 performance can be improved by structuring the HDF5 data file in a “chunked” format, depending on the patterns of data access in actual use. A chunk is a contiguous two-dimensional block of the data array that is stored and retrieved as a unit whenever any element within it is accessed. The results reported above were obtained without chunking. Separate tests were performed to compare the

unchunked format with four different (marker, sample) chunk dimensions: (256, 256), (1000, 1000), (1000, 5258), and (1, 5258). Loading the HDF5 file was no faster for any of these configurations, and 7-fold slower for (1000, 1000). Retrieval times were also not improved by chunking, using the *H5Sselect_hypherslab* procedure described above.

Spark Implementation

Spark is an open source distributed general purpose cluster-computing framework. Spark version 2.3.0 was used, using the PySpark (Python 3.6.5) interface. The Java VM used was Java Hotspot 64 bit 1.8.0_171. It is possible there might be a small improvement in benchmark results using the pure Scala interface for Spark, though in general the overhead for using PySpark is not large. Spark benchmarks were run by reading from pre-prepared Parquet file-format version of the genotype matrix. Parquet is an open source file format for Hadoop and stores nested data structures in a flat columnar format. Besides Spark, Parquet can be used in any Hadoop ecosystem like Hive, Impala and Pig. Data was stored in sample-fast orientation with samples in columns and markers in rows, and vice versa for marker-fast orientation.


MonetDB Implementation

MonetDB is an open source column-oriented database management system, and version 1.1 of MonetDB was used for the benchmarking. Similar to Parquet file format in Spark, data was stored in sample-fast orientation with samples as columns and markers in rows. Due to number of column restriction in MonetDB, we were not able to store data in marker-fast orientation as the number of markers exceeded the limit for number of columns. MonetDB out of the box is configured to use any available memory on the server.

3 Results and Discussion

3.1 Data Loading

The load times for each system are presented in Figure 2. HDF5 was the fastest, with MongoDB also performing reasonably well. The two RDMS performed poorly, with MariaDB being the worst, taking approximately 90 times longer than HDF5. While loading time is a lower priority than extraction times, the tight turnaround times from receiving marker data from the lab and generating genomic predictions for selection makes loading times of more than a day for large datasets undesirable for routine implementation of GS. While the process used for loading the data was not optimized, it is unlikely that loading times for MariaDB could be reduced to an acceptable level of performance. The performance of PostgreSQL could potentially be reduced to less than 1 day, but the large gap in performance compared to HDF5, MonetDB, and MongoDB is undesirable. Based on some preliminary testing (data not reported) MongoDB outperformed Elasticsearch (ES). To achieve optimal performance, ES settings must be tuned according to each dataset, hardware environment and expected response time. Furthermore, ES is not designed to return large amounts of data at once, using a scrolling API instead, which complicates the task of gathering query results. Based on these initial results and the complexity of implementing ES as part of GOBii-GDM solution, the decision was made to not pursue further benchmarking on ES.



3.2 Data Extraction

Data extraction tests intentionally extended to very large result sets to test performance beyond the limits expected in actual use. Figure 3 shows the extraction times for increasing number of markers, for all samples in the dataset. For a contiguous block of markers, HDF5 showed the best performance, with the next best solution, MonetDB, performing 11 times slower. On the other hand, for a random list of markers, although HDF5 shows better performance overall, Spark showed a steady performance across different marker blocks, and seems to outperform HDF5 at high marker numbers. The big discrepancy in performance of Spark between contiguous and random list of markers can be explained in the columnar nature of Spark. Spark Parquet file format is a column-oriented data format, so it does not have an "index" of rows and their order, so to ask for a contiguous chunk of rows is antithetical to the design of the data format. Results for MariaDB were greater than 25 hours for all points, off scale in Figure 2a and 2b. When selecting random markers for all samples, MonetDB extraction times exceeded 25 hours for even modest numbers of markers. This is likely due to the orientation of the data stored in the system. Due to limitations in the number of columns for the MonetDB and Spark set-up used for benchmarking, data could not be stored with markers as the columns (> 31 million columns). Given MonetDB is designed for fast querying of columns, the system did not perform well for marker (row) queries.

Figure 2: Load times for database systems

Benchmarking database systems for GS

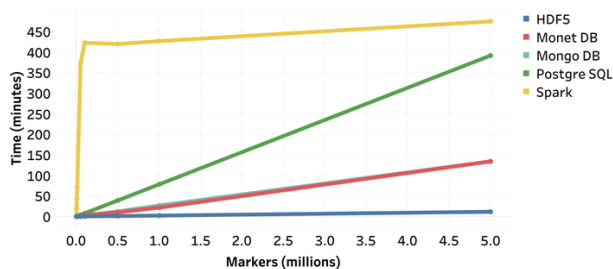


Figure 3a: Times for extracting set of contiguous markers for all samples. Times for MongoDB and MonetDB where essentially identical




Figure 3b: Times for extracting set of random markers for all samples

Extraction times for retrieving all markers for subsets of samples are shown in Figure 3. Again, the results for MariaDB were greater than 25 hours for all queries. As expected MonetDB performed significantly better with queries on samples (columns). As with queries on markers, HDF5 performed best, with the next best solution, Spark, 1.2 and 2.2 times slower than HDF5 for the contiguous and random sample scenarios respectively. For queries on samples, the relative performance of PostgreSQL dropped substantially, with extract times exceeding 20 hours for all data points. The drop in relative performance, and insensitivity to number of samples in the random sample scenario, may be related to the way in which the data is stored in PostgreSQL. A JSONB object is stored for each marker, with each object containing the allele calls for all samples. The observed performance indicates that the total extraction time is influenced more strongly by the time to fetch the JSONB objects from disk into memory, than to extract data for desired samples from the JSONB objects in memory. For the queries extracting all markers from selected samples, two data orientations were tested for MongoDB, with the sample-fast orientation (shown) performing significantly better.




Figure 4a: Times for extracting set of contiguous samples for all 32 million markers




Figure 4b: Times for extracting set of random samples for all 32 million markers

Results from USECASE III, extraction based on varying lists of markers and samples. Results for a list of 1 million markers and varying numbers of samples is shown in Figure 5. Once again HDF5 gave the best performance with the next best system performing 5.7 and 1.8 times slower for the largest contiguous and random marker and sample lists, respectively. For the contiguous scenario, MonetDB gave the second-best performance, twice as fast as MongoDB, but exceeded 48 hours for the random list selecting data from 4000 samples and 1 million markers. Neither of the RDBMS performed well for scenario 3, with all extractions taking more than 48 hours for 4000 samples and 1 million markers.




Figure 5a: Times for extracting a cross-section of a number of contiguous samples and 1 million contiguous markers




Figure 5b: Times for extracting a cross-section of a number of random samples and 1 million random markers

4 Conclusion

For all USECASES the orientation of the data storage had a substantial impact on the performance time of extraction. The fact that PostgreSQL and MonetDB had limitations on storing the data in sample fast and marker fast orientations, reduces their utility for systems that would be regularly queried based on either markers or samples. For systems using HDF5 or MongoDB, best performance would be obtained by storing the

data in both orientations with queries being directed to the optimal orientation. While HDF5 showed consistently superior performance in extraction times and loading, implementation in a genomic data management system would require a hybrid approach, with critical meta information likely stored in a RDBMS. A final determination on whether to build a system around HDF5 would need to account for the performance and complexity of developing and deploying a hybrid system. All performance tests were done using a fixed number of cores, but previous studies have shown that the performance of NoSQL distributed file systems, such as MongoDB and Spark, increase with access to more cores [8]. Further testing is required to determine if the performance of MongoDB or Spark would surpass HDF5 when deployed on large clusters.

Funding

This work has been supported by the Bill & Melinda Gates Foundation and Agropolis Foundation grant E-SPACE (1504-004)

Conflict of Interest: none declared.

References

- [1] T. H. E. Meuwissen, B. J. Hayes, and M. E. Goddard, "Prediction of Total Genetic Value Using Genome-Wide Dense Marker Maps," *Genetics*, vol. 157, no. 4, p. 1819 LP-1829, Apr. 2001.
- [2] J. M. Hickey, T. Chirugwi, I. Mackay, W. Powell, and I. G. S. in C. B. P. W. Participants, "Genomic prediction unifies animal and plant breeding programs to form platforms for biological discovery," *Nat. Genet.*, vol. 49, p. 1297, Aug. 2017.
- [3] Z. Lin, B. J. Hayes, and H. D. Daetwyler, "Genomic selection in crops, trees and forages: a review," *Crop Pasture Sci.*, vol. 65, no. 11, pp. 1177–1191, 2014.
- [4] S. Wang et al., "High dimensional biological data retrieval optimization with NoSQL technology," in *BMC genomics*, 2014, vol. 15, no. 8, p. S3.
- [5] U. Röhm and J. Blakeley, "Data management for high-throughput genomics," *arXiv Prepr. arXiv:0909.1764*, 2009.
- [6] M. M. Hoffman, O. J. Buske, and W. S. Noble, "The Genomdata format for storing large-scale functional genomics data," *Bioinformatics*, vol. 26, no. 11, pp. 1458–1459, 2010.
- [7] N. Alexandrov et al., "SNP-Seek database of SNPs derived from 3000 rice genomes," *Nucleic Acids Res.*, vol. 43, no. D1, pp. D1023–D1027, 2014.
- [8] R. Cijvat et al., "Genome sequence analysis with MonetDB," *Datenbank-Spektrum*, vol. 15, no. 3, pp. 185–191, 2015.
- [9] V. Guimaraes et al., "A study of genomic data provenance in NoSQL document-oriented database systems," in *Bioinformatics and Biomedicine (BIBM), 2015 IEEE International Conference on*, 2015, pp. 1525–1531.
- [10] G. Manyam, M. A. Payton, J. A. Roth, L. V. Abruzzo, and K. R. Coombes, "Relax with CouchDB—Into the non-relational DBMS era of bioinformatics," *Genomics*, vol. 100, no. 1, pp. 1–7, 2012.
- [11] M. Gabetta, I. Limongelli, E. Rizzo, A. Riva, D. Segagni, and R. Bellazzi, "BigQ: a NoSQL based framework to handle genomic variants in i2b2," *BMC Bioinformatics*, vol. 16, no. 1, p. 415, 2015.
- [12] W. L. Schulz, B. G. Nelson, D. K. Felker, T. J. S. Durant, and R. Torres, "Evaluation of relational and NoSQL database architectures to manage genomic annotations," *J. Biomed. Inform.*, vol. 64, pp. 288–295, 2016.
- [13] E. Dede, M. Govindaraju, D. Gunter, R. S. Canon, and L. Ramakrishnan, "Performance evaluation of a mongodb and hadoop platform for scientific data analysis," in *Proceedings of the 4th ACM workshop on Scientific cloud computing*, 2013, pp. 13–20.
- [14] G. Sempéré, F. Philippe, A. Dereeper, M. Ruiz, G. Sarah, and P. Larmande, "Gigwa—Genotype investigator for genome-wide analyses," *Gigascience*, vol. 5, no. 1, p. 25, 2016.
- [15] M. D. McMullen et al., "Genetic properties of the maize nested association mapping population," *Science (80-.)*, vol. 325, no. 5941, pp. 737–740, 2009.
- [16] J. C. Glaubitz et al., "TASSEL-GBS: a high capacity genotyping by sequencing analysis pipeline," *PLoS One*, vol. 9, no. 2, p. e90346, 2014.
- [17] C. Gormley and Z. Tong, *Elasticsearch: The Definitive Guide: A Distributed Real-Time Search and Analytics Engine*. " O'Reilly Media, Inc.," 2015.
- [18] M. Zaharia et al., "Apache spark: a unified engine for big data processing," *Commun. ACM*, vol. 59, no. 11, pp. 56–65, 2016.

- [19] Robert Bukowski, Xiaosen Guo, Yanli Lu, Cheng Zou, Bing He, Zhengqin Rong, Bo Wang, Dawen Xu, Bicheng Yang, Chuanxiao Xie, Longjiang Fan, Shibin Gao, Xun Xu, Gengyun Zhang, Yingrui Li, Yiping Jiao, John F Doebley, Jeffrey Ross-Ibarra, Anne Lorant, Vince Buffalo, M Cinta Romay, Edward S Buckler, Doreen Ware, Jinsheng Lai, Qi Sun, Yunbi Xu; Construction of the third-generation *Zea mays* haplotype map, *GigaScience*, Volume 7, Issue 4, 1 April 2018, gix134, <https://doi.org/10.1093/gigascience/gix134>
- [20] Mason C.E. et al. (2010) Standardizing the Next Generation of Bioinformatics Software Development with BioHDF (HDF5). In: Arabnia H. (eds) *Advances in Computational Biology. Advances in Experimental Medicine and Biology*, vol 680. Springer, New York, NY
- [21] Nicholas Loman, Aaron Quinlan , Porettools: a toolkit for analyzing nanopore sequence data, *bioRxiv* 007401; doi: <https://doi.org/10.1101/007401>

- [22] Thomson, Michael J. "High-Throughput SNP Genotyping to Accelerate Crop Improvement." (2014).

Appendix 1

Table 1: Results for extracting M markers for all samples

Contiguous	# of markers (mil)	MonetDB	HDF5	MongoDB	PostgreSQL	MariaDB	Spark
Yes	0	0	0	0	0	0	0
Yes	0.001	0.1	0	0	0.1	1566	19.5
Yes	0.005	0.2	0	0.1	0.5	1623	43.45
Yes	0.01	0.3	0	0.3	0.8	1523	71.3
Yes	0.05	1.2	0.1	1.4	3.9	1579	373.08
Yes	0.1	3.7	0.3	2.5	7.8	1663	423.07
Yes	0.5	10	1.2	12.3	39.2	1520	420.32
Yes	1	21.7	2.3	26.8	78.4	1632	426.8
Yes	5	134.2	11.6	134.3	391.8		475
No	0	0	0	0	0	0	0
No	0.001	44.5	0.1	0.2	0.5	1545	252
No	0.005	47.2	0.4	0.7	0.6	2043	251
No	0.01	47.8	0.9	1.8	1.2	2108	251
No	0.05	58.6	4.4	7.1	6	3017	252
No	0.1	74.8	8.6	13.7	13.2	6021	252
No	1	2700	80.8	130.8	149.5		264
No	5		443	858.8	945.6		316

Table 2: Results for extracting N samples for all markers

Contiguous	# of samples	MonetDB	HDF5	MongoDB	PostgreSQL	MariaDB	Spark
Yes	0	0	0	0	0	0	0
Yes	50	1.8	2	190.8	1282	1830	2.7
Yes	100	4	4	195.8	1301	1868	2.82
Yes	500	42.1	19.6	234.4	1401		13.73
Yes	1000	103.2	40.5	291.4	1528		32.63
Yes	2000	238.5	84.7	405.5	1804		92.52
			128.				149.3
Yes	3000	432.5	4	518.2	2069		8
							211.5
Yes	4000	698.5	169	632.8	2327		2
No	0	0	0	0	0	0	0

Benchmarking database systems for GS

No	50	1.9	1.7	188.5	1482	1575.63	2.9
No	100	4.7	2.4	195.6	1505	1535.53	3.3
No	500	48.7	15.6	239.3	1648		14
No	1000	118.3	22.6	295.7	1696		35.5
No	2000	306.6	58.5	409	1678		108
No	3000	547	81.6	621.3	1674		171
No	4000	626.9	107.7	677.9	1681		239

Table 3: Results for extracting M markers by N samples

Contiguous?	# of markers	# of Samples	MonetDB	HDF5	MongoDB	PostgreSQL	MariaDB	Spark
Yes	1k	0	0.00	0.0	0.00	0.00	0.00	0.00
Yes	1k	1000	0.00	0.0	0.00	0.10	1287	1.60
Yes	1k	3000	0.10	0.0	0.00	0.10	1684	8.55
Yes	1k	4000	0.10	0.0	0.00	0.10	1718	13.37
No	1k	0	0.00	0.0	0.00	0.00	0.00	0.00
No	1k	1000	8.40	0.1	0.20	0.20	1073	20.4
No	1k	3000	28.60	0.1	0.20	0.20	1163	144
No	1k	4000	24.70	0.1	0.20	27.10	1662	201
Yes	50k	0	0.00	0.00	0.00	0.00	0.00	0.00
Yes	50k	1000	0.10	0.10	0.70	2.60	1726	38.08
Yes	50k	3000	0.40	0.10	1.20	3.40	1635	204.3
Yes	50k	4000	0.80	0.20	1.50	3.80	1111	299.3
No	50k	0	0.00	0.00	0.00	0.00	0.00	0.00
No	50k	1000	10.80	4.10	6.60	5.00	1297	20.05
No	50k	3000	33.20	4.20	7.20	5.20	1623	145
No	50k	4000	32.40	4.20	8.00	4.80	1756	203
Yes	1m	0	0.00	0.00	0.00	0.00	0.00	0.00
Yes	1m	1000	3.40	0.90	12.10	49.90	1045.00	41.32
Yes	1m	3000	10.60	2.10	22.10	68.20	1048.00	2
Yes	1m	4000	14.70	2.60	29.60	76.10	1492.00	314.2
No	1m	0	0.00	0.00	0.00	0.00	0.00	0.00
No	1m	1000	192.1	72.6	109.60	121.80	1802.00	22.77
No	1m	3000	1846.6	78.5	138.10	143.40	1911.00	152.5
No	1m	4000	3291.7	83.0	153.90	151.00	1935.00	212.6