

Transparent Incremental Updates for Genomics Data Analysis Pipelines

Edvard Pedersen¹, Nils Peder Willassen², Lars Ailo Bongo¹

¹Department of Computer Science, University of Tromsø, Norway

²Department of Chemistry, University of Tromsø, Norway

^{1&2}Centre for Bioinformatics, University of Tromsø, Norway

edvard.pedersen@uit.no, nils-peder.willassen@uit.no,
larsab@cs.uit.no

Abstract. A large up-to-date compendium of integrated genomic data is often required for biological data analysis. The compendium can be tens of terabytes in size, and must often be frequently updated with new experimental or meta-data. Manual compendium update is cumbersome, requires a lot of unnecessary computation, and it may result in errors or inconsistencies in the compendium. We propose a transparent file based approach for adding incremental update capabilities to unmodified genomics data analysis tools and pipeline workflow managers. This approach is implemented in the GeStore system. We evaluate GeStore using a real world genomics compendium. Our results show that it is easy to add incremental updates to genomics data processing pipelines, and that incremental updates can reduce the computation time such that it becomes practical to maintain large-scale up-to-date genomics compendia on small clusters.

1 Introduction

Recent advances in scientific instruments, such as next-generation sequencing machines, has the potential of producing data that provide views of biological processes at different resolutions and conditions, opening a new era in molecular biology and molecular medicine [1]. Many of the data analysis techniques developed for analyzing such genomic data integrate data from many experiments with metadata from multiple knowledge bases. The information in the knowledge bases [2] is essential for understanding the biological content of the experiment data. For example, the results of DNA sequencing may not become truly useful before the UniProt [3] knowledge base is used to map sequence bases to genes, the per gene results are compared to results from other experiments, and the significant differences have been mapped to biological functions using the Go [4] knowledge base.

Genomic data integration and analysis is typically implemented as a pipeline with several tools, where the output files of one tool acts as the input files for the next tool. The specific set of tools to use depends on the biological problem that is being investigated. Often large amounts of data must be analyzed, since new sequencing machines produce multiple terabytes of data per experiment [5]. The cost of the analysis can

therefore be orders of magnitude larger than the cost of creating the data to be analyzed [6]. Cost efficient data analysis is therefore a key challenge for genomics data analysis

Integrating new experimental data or updating meta-data may provide novel biological insights. It is therefore important to update a biological compendium when new data becomes available. Meta-data updates are especially important since it represents the state of knowledge in the field [7]. However, current biological analysis tools typically require recalculating the entire compendium for meta-data updates. Such full updates increase the computational time and cost; often to the point where reanalysis is not done.

The cost of reanalysis can be greatly reduced by using incremental updates [8] that limit recomputation to new and updated data. We believe such an approach for incremental updates of genomic data must satisfy the following four requirements. First, most existing analysis tools should be supported without any modifications since it is not practical to maintain modified versions of the many analysis tools used in genomic data analysis pipelines. Transparent incremental updates are therefore needed. Second, the incremental updates should be independent of the job and resource management systems used to run the pipeline tools since genomic analysis pipelines are run on many different platforms. Third, it should be easy to add incremental update support to an existing pipeline. The system should therefore handle update detection, processing of incremental updates, and integration of the incremental update with a previous full update. Fourth, it should scale to large-scale compendium.

To our knowledge, no previous incremental update systems for large-scale data [9–15] satisfy all four requirements. These provide the required functionality and scalability, but do not provide the easy to use transparent incremental updates that are necessary to add incremental updates to existing pipelines. Instead they require either porting applications to a specific framework such as Dryad [16] or MapReduce [17], or writing ad-hoc scripts for input generation and output merging.

We present the GeStore system for incremental update management. GeStore uses a transparent file based approach that satisfies all four requirements. Most pipeline tools take as input one or more files with input and meta-data, and produce output consisting of one or more files. Incremental updates can therefore be implemented by modifying the input or meta-data files such that these only contain the data for an incremental update, and then merging the incremental output with the previous results. Tool code is unmodified, and the only modifications to the pipeline are two GeStore calls for generating input files and merging output files. GeStore provides a plugin framework for implementing parsers, tool-specific incremental file generators, and tool-specific output file mergers. GeStore uses the Hadoop software stack for scalable data processing.

Our contributions are threefold: (i) we propose a promising approach for adding incremental updates to unmodified genomic data analysis pipelines, leading to substantial reduction in time and resources needed to update large biological compendium, (ii) we present the design and implementation of our system, including a framework for implementing plugins that enable transparent incremental updates, and (iii) we present the feasibility of our approach and initial experimental evaluation of our system using a real metagenomics analysis pipeline and real data.

2 GeStore

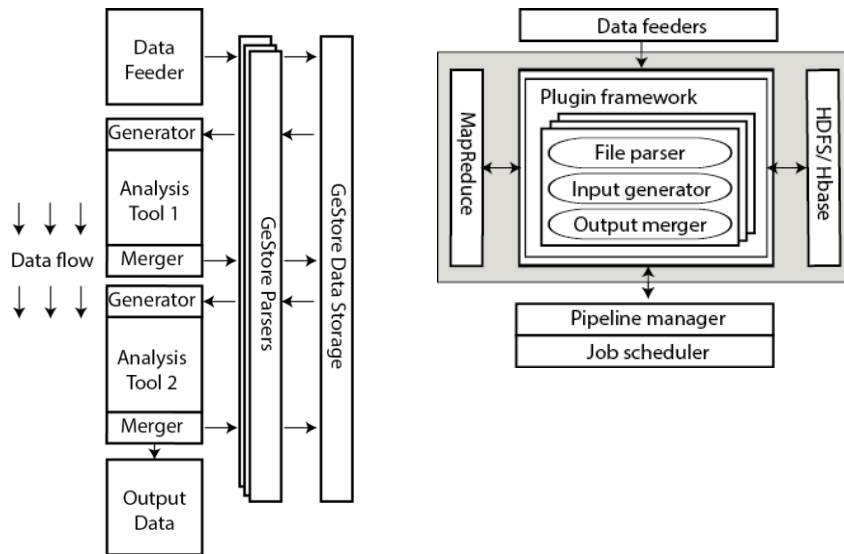


Fig. 1. Incremental pipeline execution (left). GeStore architecture (right).

GeStore is a system for enabling transparent incremental computations for unmodified file-based data analysis pipelines. GeStore consists of a runtime system that provides a plugin framework for incremental input file generation and output file merging, a tool-set for parsing and detecting changes in files, and data storage and management (Fig. 1, left). GeStore exports an interface for data feeders, and an interface for workflow managers. Data feeders are typically scripts that periodically download new input data or updated meta-data from remote repositories or local storage systems. All downloaded files are stored in GeStore. Pipeline managers (or pipeline configurations) are modified to call GeStore before running each tool in the pipeline in order to generate the incremental input files used by the tools and to merge the resulting files (Fig. 1, right). GeStore uses HDFS [18] and Hbase [19] to efficiently store incremental files and GeStore meta-data, and Hadoop MapReduce [17] to run scalable change detection jobs. In addition GeStore comprise library functions and tools to add incremental updates to pipeline tools, and client applications to access data stored within GeStore.

2.1 File based incremental updates

GeStore uses a transparent file based approach where incremental updates are implemented by modifying the input and meta-data files read by genomics data analysis tools such that these only contain the data for incremental update computations. The tool will then be run as normal, but it will typically produce a partial result. The partial result is merged with previously produced results and stored in GeStore.

We have chosen a file based approach since there are relatively few file formats that are used by many genomics applications. It is therefore feasible to implement parsers that support most file formats and therefore most genomics pipeline tools. In addition, most file formats are simple and structured which makes it easy to write parsers for each format. However, update file generation and output merging is not trivial to implement. It may be necessary to understand the biological content of the data and how the tools read and write the data. For example, for the widely used BLAST [20] tool most changes to the UniProt [3] input data records are for fields that are not used in the computation, and the output data records contains a field (e-value [15]) that is incorrect for incremental updates. Both of these issues can be fixed by writing relatively simple code for ignoring the insignificant fields during change detection and by fixing the e-values in the output data. In addition, the system must provide, low overhead storage for incremental update data, and efficient generation of incremental update files.

GeStore provides an interface that the pipeline system uses to request one or more incremental update input files, and to merge the partial results with previously produced results. These upcalls can be added by adding stages to the pipeline configuration before and after the execute tool step. Alternatively, the GeStore calls can be added by modifying the code in the pipeline manger that manages the lists of input, meta-data, and output files used to setup a tool for execution and to store the results.

2.2 Plugin framework

GeStore provides a plugin framework to support many different file formats and pipeline tools. To add incremental updates to a pipeline the administrator must first write a plugin for each tool in the pipeline. These are then used by GeStore for each incremental update. A plugin comprise three parts: (i) a parser for each file type used by the pipeline tool, (ii) tool-specific incremental file generator, and (iii) tool-specific incremental output file merger. Each plugin has typically a few tens of lines of code. Many plugins also share parsers and file mergers, and GeStore provides many library functions for parsing, change detection, and merging of files. The plugins are managed by a framework that provides efficient data storage, and low overhead file parsing, generation, and merging.

The file parser must define schemas for the input files and meta files used by a tool, and implement six methods that: (i) provide regular expressions that define the start and end of an entry in the file, (ii) split an entry into columns, (iii) compare two versions of an entry, (iv) check if an entry is well-formed, (v) put the entry into HBase, and (vi) generate output in other formats. The file generator requires implementing one method that specifies the parsers to use for each file format, and the fields to write to the input file. The output merger requires implementing a method to merge the output data with previously produced output data stored in GeStore. This may include minor fixes to output data fields as discussed above.

2.3 Data storage and management

GeStore uses the Hadoop software stack for scalable data storage and processing. GeStore maintains: (i) HBase tables and HDFS files with pipeline tool input, output and meta-data file data, (ii) a cache of previously generated incremental update files stored in HDFS, and (iii) HBase tables with per plugin instance state that is used to generate file, merge output files, and provide provenance information.

For file types that have a parser implemented, the data is split into entries and entry fields. These are stored as rows and columns in HBase using a file-format specific schema. The only required column in the schema is a unique ID for each row. The HBase schema can be modified by adding new columns to the table if for example the file format changes or the parser is modified. GeStore use the versioning mechanism in HBase to store only updated fields (i.e. delta compression), and to return the data for a given time period. The timestamp for a version corresponds to either the file generation date, release date, or version.

3 Incremental METApipeline

METApipeline is used by our biology collaborators to find novel enzymes by analyzing sequence data from marine microbial communities. METApipeline is currently run using the GePan pipeline management system (developed by Tim Kahlke at the University of Tromsø). It includes the following tools:

1. Multiple Genome Aligner (MGA): [21] does alignment of closely related DNA sequences. It does not require meta-data from knowledge bases.
2. MGA-exporter: converts the MGA output to the format used by the next stage.
3. FileScheduler: partitions and distributes the input data to the compute nodes.
4. Protein BLAST (BLASTP) [20]: maps sequences to information from the UniProt Swiss-Prot and TrEMBL [3] knowledge bases.
5. HMMer [22]: maps sequences to information from the Pfam-A and Pfam-B [23] protein family databases.
6. Annotator: gathers the results from the preceding tools, and converts the data to a custom format.
7. Annotator-exporter: converts the annotator output to a format that can be used by data visualization and exploration tools.

To add incremental updates to METApipeline we had to write parsers for the six file formats used by the pipeline: FASTA, UniprotKB meta-data, Pfam meta-data, BLAST output, HMMer output, and MGA output. We also had to write plugins for the BLAST and HMMer tools. The BLAST plugin corrects incremental e-values as discussed in [15] during merge. The HMMer plugin only generates input files.

The file format plugins were a total of 844 lines of Java code, and the tool plugins were 283 lines of Java code. The results show that file based incremental could be used for all tools in METApipeline, and there are relatively few lines of tool specific code.

To integrate GeStore with METApipe, we modified the code that generates the Grid Engine [24] scripts that run the pipeline tool code. GeStore calls were added by changing the file initialization commands to GeStore calls. In total, about 120 lines of code were changed in METApipe. We expect the changes required to other pipeline management systems to be similarly small.

4 Evaluation

Our initial experimental evaluation compares the benefits and overheads of using incremental updates for the METApipe metagenomics analysis pipeline. Our experiments were run on a small cluster with one frontend and eight compute nodes. Each node is equipped with two Intel Xeon E5-1620 CPUs running at 3.6 GHz and 32 gigabytes of RAM. They have a total of 4.5 TB of local HDD. They also have 2.6 TB of NFS storage shared between them. The machines are connected using gigabit Ethernet. We believe such a cluster configuration is realistic for research labs that maintain genomic compendia.

We use a 15 mega base pairs metagenomics dataset from the Yellowstone Park [25] as input data. Processing this small dataset takes 2.5 hours on our small cluster. We incrementally update the dataset on the last day of the month from January 2011 to July 2011. There were 6 updates to Uniprot Swiss-Prot and TrEMBL, and one update to Pfam A and B in that period.

4.1 Update relationships

| | Total entries | Total updates | Significant | New entries |
|------------|---------------|---------------|-------------|-------------|
| Swiss-Prot | 527590 | 38.76% | 0.44% | 0.40% |
| TrEMBL | 14738346 | 32.11% | 4.89% | 4.88% |
| Pfam-A | 1076 | 100.00% | 100.00% | 3.25% |

Table 1. Monthly meta-data collection updates between January-July 2011. Averages reported.

To analyze relationships between meta-data changes and input file changes, we averaged all changes in UniProt TrEMBL, UniProt Swiss-Prot, and Pfam-A meta-data collections between January and July in 2011 (Table 1). In Swiss-Prot and TrEMBL most changes are to annotation that does not require BLAST recomputation, and hence a significant difference in incremental update execution time. Pfam has a naïve plugin that marks all changes as significant, and has therefore a high rate of significant changes (100%). The Pfam plugin could be improved by doing more precise classifications of non-significant updates. These results demonstrate the benefits of tool specific plugins.

4.2 GeStore Improvements and overhead

We measured METApipe execution time for full updates and incremental updates with 1, 3, and 6 month periods (Table 2). The analysis time is dominated by BLASTp. Since

BLAST execution time scales linearly with the input size, the smaller incremental input data generated by GeStore significantly reduce BLAST analysis time, and hence total execution time.

| | Analysis | Overhead | Total |
|---------------------------|----------|----------|-------|
| Full update (Jan 2011) | 9141 | 0 | 9141 |
| with GeStore | 10718 | 2562 | 13280 |
| Incremental (Jan – Feb) | 893 | 755 | 1647 |
| Incremental (Jan – April) | 1736 | 3497 | 5233 |
| Incremental (Jan – Juli) | 2850 | 3736 | 6586 |

Table 2. METApipe execution time split into analysis time and GeStore overhead (all in seconds).

GeStore has an overhead for HMMer of 800 seconds when generating a complete database, and 300 seconds when retrieving a cached database. Generating an incremental update database takes 2800 seconds, this is because the Pfam plugin marks all updates as important. BLASTp has an overhead of 1700 seconds for generating a full database. The incremental update time is 300 to 800 seconds depending on the size of the update.

Although GeStore overhead is significant for these experiments it will be much smaller for bigger, more realistic, input dataset sizes since the analysis time depends on input data size, while GeStore overhead depends on meta-data size. In addition, we expect to reduce the Pfam change detection overhead by implementing data aware change detection (as discussed above).

The storage overhead increases sub linearly for UniProt since there are relatively few updates per month (as shown above), the January UniProt database file size is 33 Gb. When stored in HBase it requires 48 Gb of space. However, the total size of the UniProt databases is 252 Gb, but only requires 77 Gb of space in GeStore. For Pfam the storage requirements increase linearly, from to 3.3 Gb to 7.1 Gb and 2.9 Gb and 6.3 Gb respectively for GeStore and total file size. The storage requirements can be significantly reduced by improving the plugin for HMMer.

GeStore achieves similar analysis runtime improvements (90%, for 5% meta-data updates) to incremental BLAST as reported in [15]. Execution time improvements ranging from 20% to 99% are reported in [9–14], but for applications from the data center domain. We have not experimentally compared the execution time improvements and overheads to other large scale incremental update tools since these require modifications to the pipeline tools.

5 Related Work

Systems and frameworks for incremental updates on large scale datasets include Incoop [10], Percolator [11], Nectar [9], DryadInc [12], CBP [13], and HaLoop [14]. In Percolator and CBP the programmer implements a system specific incremental program using respectively event-driven mini transactions and stateful primitives. Incoop, Nectar,

DyradInc, and HaLoop use data dependency graphs of Dryad [16] or MapReduce programs to automatically replace the input data for a computation with previously calculated results. GeStore combines these two main approaches; a programmer implements file generators and mergers for unmodified programs. GeStore is independent of the programming model and job management system, so the applications can be executed using Dryad, MapReduce [11], or the Grid Engine [24].

GeStore extends the work in [15] by providing a framework and libraries to implement the necessary pre and post processing of data moved between a data warehouse and genomic analysis tools. This makes it easier to add additional support for additional genomic analysis tools as we have demonstrated by implementing incremental updates for a complete metagenomics analysis pipeline.

Simple change detection is supported by tools such as Unix diff, delta encoding compression systems [26], and version management systems such as CVS [27]. However, the change detection in these do not take into account the complex inter-file relationships found in genomic datasets.

The file tables maintained by GeStore are similar to declarative views maintained by data warehouses [28]. Incremental updates have also been used for non-distributed computation result caching (memoization) as in [8].

Popular approaches for genomics pipeline management are Galaxy [29] and Bio-Conductor [30]. These do not provide incremental computation.

We evaluated GeStore using the locally developed METApipe pipeline. An alternative is the JCVI metagenomics analysis pipeline [31].

6 Conclusions and Future work

We proposed an approach for adding incremental updates to unmodified genomic data analysis pipelines, leading to substantial reduction in time and resources needed to update large biological compendiums. We presented the design and implementation of the GeStore system, including a framework for implementing plugins that enable transparent incremental updates. We demonstrated the feasibility of our approach and provided an initial experimental evaluation of our system using a real metagenomics analysis pipeline and real data. The cost effective transparent incremental updates provided by GeStore makes it practical to frequently update large genomic compendium with new experimental and meta-data, and thereby enabling novel biological discoveries.

We plan to further evaluate the benefits and overhead of incremental updates for genomics data analysis by applying GeStore to the pipeline producing data for the IMP [32] tool, and to a Galaxy [29] pipeline. Galaxy can also be used to provide a GUI for GeStore configuration and data management.

7 Acknowledgements

Thanks to Espen Robertsen and Tim Kahlke for help with the GePan pipeline, Jon Ivar Kristiansen for maintaining our cluster, and Martin Erntsen for his comments.

8 References

- [1] S. C. Schuster, "Next-generation sequencing transforms today's biology.," *Nature Methods*, vol. 5, no. 1, pp. 16–18, 2008.
- [2] M. Y. Galperin and X. M. Fernández-Suárez, "The 2012 Nucleic Acids Research Database Issue and the online Molecular Biology Database Collection.," *Nucleic acids research*, vol. 40, no. Database issue, pp. D1–8, Jan. 2012.
- [3] M. Magrane and U. Consortium, "UniProt Knowledgebase: a hub of integrated protein data.," *Database the journal of biological databases and curation*, vol. 2011, 2011.
- [4] M. Ashburner, C. A. Ball, J. A. Blake, D. Botstein, H. Butler, J. M. Cherry, A. P. Davis, K. Dolinski, S. S. Dwight, J. T. Eppig, M. A. Harris, D. P. Hill, L. Issel-Tarver, A. Kasarskis, S. Lewis, J. C. Matese, J. E. Richardson, M. Ringwald, G. M. Rubin, and G. Sherlock, "Gene Ontology: tool for the unification of biology.," *Nature Genetics*, vol. 25, no. 1, pp. 25–29, 2000.
- [5] S. D. Kahn, "On the Future of Genomic Data," *Science*, vol. 331, no. 6018, pp. 728–729, Feb. 2011.
- [6] J. Wilkening, A. Wilke, N. Desai, and F. Meyer, "Using clouds for metagenomics: A case study," *2009 IEEE International Conference on Cluster Computing and Workshops*, pp. 1–6, 2009.
- [7] R. Sandberg and O. Larsson, "Improved precision and accuracy for microarrays using updated probe set definitions.," *BMC bioinformatics*, vol. 8, no. 1, p. 48, Jan. 2007.
- [8] Y. A. Liu, S. D. Stoller, and T. Teitelbaum, "Static caching for incremental computation," *ACM Transactions on Programming Languages and Systems*, vol. 20, no. 3, pp. 546–585, May 1998.
- [9] P. K. Gunda, L. Ravindranath, C. A. Thekkath, Y. Yu, and L. Zhuang, "Nectar: automatic management of data and computation in datacenters," in *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, 2010, pp. 1–8.
- [10] P. Bhatotia, A. Wieder, R. Rodrigues, U. A. Acar, and R. Pasquini, "Incoop : MapReduce for Incremental Computations," in *Proceedings of the 2nd ACM Symposium on Cloud Computing*, 2011, p. 7.
- [11] D. Peng and F. Dabek, "Large-scale Incremental Processing Using Distributed Transactions and Notifications," in *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, 2010, vol. 2006, pp. 1–15.
- [12] L. Popa, M. Budiui, Y. Yu, and M. Isard, "DryadInc: reusing work in large-scale computations," *Proceedings of the 2009 conference on Hot topics in cloud computing*, p. 21, Jun. 2009.
- [13] D. Logothetis, C. Olston, B. Reed, K. C. Webb, and K. Yocum, "Stateful bulk processing for incremental analytics," in *Proceedings of the 1st ACM symposium on Cloud computing - SoCC '10*, 2010, p. 51.
- [14] Y. Bu, B. Howe, M. Balazinska, and M. D. Ernst, "The HaLoop approach to large-scale iterative data analysis," *The VLDB Journal*, vol. 21, no. 2, pp. 169–190, Mar. 2012.
- [15] G. Turcu, S. Nestorov, and I. Foster, "Efficient Incremental Maintenance of Derived Relations and BLAST Computations in Bioinformatics Data Warehouses," in *Data Warehousing and Knowledge Discovery*, vol. 5182, I.-Y. Song, J. Eder, and T. Nguyen, Eds. Springer, 2008, pp. 135–145.
- [16] M. Isard, M. Budiui, Y. Yu, A. Birrell, and D. Fetterly, "Dryad: distributed data-parallel programs from sequential building blocks," *ACM SIGOPS Operating Systems Review*, vol. 41, no. 3, p. 59, 2007.
- [17] J. Dean and S. Ghemawat, "MapReduce: a flexible data processing tool," *Communications of the ACM*, vol. 53, no. 1, p. 72, Jan. 2010.

- [18] K. Shvachko, H. Kuang, S. Radia, and R. Chansler, "The Hadoop Distributed File System," *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies MSST*, vol. 0, no. 5, pp. 1–10, 2010.
- [19] Apache, "Apache HBase," 2012. [Online]. Available: <http://hbase.apache.org/>. [Accessed: 24-Apr-2012].
- [20] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman, "Basic local alignment search tool," *Journal of Molecular Biology*, vol. 215, no. 3, pp. 403–410, 1990.
- [21] M. Höhl, S. Kurtz, and E. Ohlebusch, "Efficient multiple genome alignment.," *Bioinformatics*, vol. 18 Suppl 1, no. Suppl 1, pp. S312–S320, 2002.
- [22] R. D. Finn, J. Clements, and S. R. Eddy, "HMMER web server: interactive sequence similarity searching," *Nucleic Acids Research*, vol. 39, no. Web Server issue, pp. W29–W37, 2011.
- [23] R. D. Finn, J. Mistry, J. Tate, P. Coghill, A. Heger, J. E. Pollington, O. L. Gavin, P. Gunasekaran, G. Ceric, K. Forslund, L. Holm, E. L. L. Sonnhammer, S. R. Eddy, and A. Bateman, "The Pfam protein families database," *Nucleic Acids Research*, vol. 38, no. Database issue, pp. D211–D222, 2010.
- [24] "Oracle Grid Engine," <http://www.oracle.com/us/products/tools/oracle-grid-engine-075549.html>. [Online]. Available: <http://www.oracle.com/us/products/tools/oracle-grid-engine-075549.html>. [Accessed: 02-May-2012].
- [25] D. Bhaya, A. R. Grossman, A.-S. Steunou, N. Khuri, F. M. Cohan, N. Hamamura, M. C. Melendrez, M. M. Bateson, D. M. Ward, and J. F. Heidelberg, "Population level functional diversity in a microbial community revealed by comparative genomic and metagenomic analyses.," *The ISME journal*, vol. 1, no. 8, pp. 703–713, 2007.
- [26] F. Douglis and A. Iyengar, "Application-specific Delta-encoding via Resemblance Detection," in *Proceedings of the USENIX Annual Technical Conference*, 2003, pp. 113–126.
- [27] D. Grune, "Concurrent Versions System, A Method for Independent Cooperation." 1986. Working paper. IR 113, Vrije Universiteit
- [28] S. Ceri and J. Widom, "Deriving Production Rules for Incremental View Maintenance," *Proceedings of the 17th International Conference on Very Large Data Bases*, pp. 577–589, Sep. 1991.
- [29] J. Goecks, A. Nekrutenko, and J. Taylor, "Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences.," *Genome biology*, vol. 11, no. 8, p. R86, Jan. 2010.
- [30] R. C. Gentleman, V. J. Carey, D. M. Bates, B. Bolstad, M. Dettling, S. Dudoit, B. Ellis, L. Gautier, Y. Ge, J. Gentry, K. Hornik, T. Hothorn, W. Huber, S. Iacus, R. Irizarry, F. Leisch, C. Li, M. Maechler, A. J. Rossini, G. Sawitzki, C. Smith, G. Smyth, L. Tierney, J. Y. H. Yang, and J. Zhang, "Bioconductor: open software development for computational biology and bioinformatics.," *Genome biology*, vol. 5, no. 10, p. R80, Jan. 2004.
- [31] D. M. Tanenbaum, J. Goll, S. Murphy, P. Kumar, N. Zafar, M. Thiagarajan, R. Madupu, T. Davidsen, L. Kagan, S. Kravitz, D. B. Rusch, and S. Yooseph, "The JCVI standard operating procedure for annotating prokaryotic metagenomic shotgun sequencing data," *Standards in genomic sciences*, vol. 2, no. 2, pp. 229–237, 2010.
- [32] A. K. Wong, C. Y. Park, C. S. Greene, L. A. Bongo, Y. Guan, and O. G. Troyanskaya, "IMP: a multi-species functional genomics portal for integration, visualization and prediction of protein functions and networks.," *Nucleic Acids Research*, vol. 40, no. Web Server issue, pp. 1–7, 2012.