

The Role of Declarative Querying in Bioinformatics

Jignesh M. Patel
University of Michigan
1301 Beal Avenue; Ann Arbor, MI 48109-2122
jignesh@eecs.umich.edu

The recent publication of a draft of the entire human genome [7, 13] has served to fuel an already explosive area of research in bioinformatics that is involved in deriving meaningful knowledge from proteins and DNA sequences [1, 4, 9-12]. Even with the full human genome sequence now in hand, scientists still face the challenges of determining exact gene locations and functions, observing interactions between proteins in complex molecular machines, and learning the structure and function of proteins, just to name a few. The progress of this scientific research is closely connected to the research in the database community in that analyzing large volumes of biological data sets involves being able to maintain and query large databases [3, 8]. Database management systems (DBMSs) could help support life sciences applications, in a number of different ways. A partial list of tasks that such applications require is: querying large structured databases (such as sequence databases), querying semi-structured (such as published manuscripts), managing data replication, querying distributed data sources, and managing parallelism in high-throughput bioinformatics. Unfortunately, current DBMSs have largely ignored supporting life sciences applications, and consequently, the life sciences researches have been forced to write tools and scripts to perform these tasks.

An interesting parallel can be drawn between the state of data management tools in life sciences, and the state of data management tools for business applications, such as a banking application, about three decades ago. Prior to the advent of the relational data model, business data was managed and queried using customized programs/scripts that were developed for each application. Reusing programs, and the algorithms for querying the data, involved rewriting application program and logic, which was very time consuming and expensive. In addition, the querying programs were closely tied to the format that was used to represent the data. Any change in the format of the data representation could break the querying programs. Furthermore writing complex queries, such as querying over multiple data sets or posing complex analytical queries, was a daunting task. One of the critical contributions of the relational data model [2] was the introduction of a declarative querying paradigm, instead of the previously used procedural paradigm. In a declarative querying paradigm, the user simply expresses the query in a high-level language, like SQL. The DBMS then figures out the best way of evaluating the query. In addition, the DBMS only presents to the user a logical view of the data against which queries are posed. The physical representation of the data can be very different from the logical view. For example, in a relational database management system (RDBMS), indices may be created, and the user doesn't have to query against the index. The user still queries against logical relations, and the system automatically determines if it is faster to use the indices to answer a query. The user is thus insulated from worrying about various details such as physical organization of data on disk, the exact location of the data, tuning the representation for better performance, choosing the best plan for evaluating a query, etc. This declarative querying paradigm has been a huge success for relational DBMSs, and today commercial RDBMSs can manage terabytes of data, and allows very complex querying on these databases. The position of this paper is that database management systems can provide similar benefits to the life sciences community, just as it did three decades ago to the business data management community. Many of the data sets that are used in life sciences are growing at an astonishing rate (such as sequence data at NCBI's GenBank [5]), and the queries that are required to process these data sets to extract biologically meaningful information are also increasing in complexity. Database management systems can and should play a significant role in managing and querying these data sets.

To better support life sciences applications, DBMSs must carefully consider the data types and the operations on the data types that are required to support these applications. One of the commonly used data type in these applications is a sequence. A sequence is often used to represent a protein's primary structure, or genes. Querying on sequences is often approximate, typically using an edit-distance matrix, such as PAM or BLOSUM matrices. Sequences have a well defined structure, and one would expect that the currently generation of commercial object-relational DBMSs, which are very good at dealing with structured data, would have support for this data type. Unfortunately, this is not the case, and such data sets are managed and queried using customized programs that run outside the DBMS. Imagine a research group that queries on such data sets frequently, and has a local copy of the sequence data set and programs to query this data set. Since new sequences are continually added to the *master* copy, keeping the data set up-to-date is now a largely manual process. If the program needs to build any auxiliary structures to process this data, such as an index [6], then this task requires customized solutions, which are cumbersome to manage. On the other hand, this is a form of replication that DBMSs are good at dealing with, and often can be managed transparently without requiring any user intervention. As another example, consider a scenario in which the user wants to query a number of different distributed sequence data sets with the same query, and then combine these results using some scoring criteria, perhaps with some post-processing on the sequence data. This task again requires customized programs, and a declarative querying interface for this task can not only make it easier to pose such queries, but can also be executed more efficiently. The reason for improved efficiency is because the DBMS can examine a number of alternative plans for evaluating the query, and can then choose the cheapest plan. Similar observations hold for other data types, such as published manuscripts and gene expression data, and DBMSs can play a significant role in managing and querying such data sets.

In summary, three decades ago DBMSs brought significant benefits to the management of business data, and today, no business application would be built without a DBMS system. This position paper argues that DBMSs can provide similar benefits in managing and querying life sciences data. Of particular interest, is the provision of a declarative querying paradigm that can be used by life sciences researchers, so that they don't have to worry about the data management, thereby allowing them to focus their time and energy on the biological sciences part of their research. As part of the Periscope project at the University of Michigan, we are examining various query language, query processing, and query optimization techniques that are required to support a declarative querying paradigm for life sciences applications.

References

1. Alberts, B. Molecular Biology of the Cell. Garland Science, New York, 2002.
2. Codd, E. F. A Relational Model of Data for Large Shared Data Banks. CACM, 13 (6). 377-387.
3. Davidson, S. B., "Tale of Two Cultures: Are There Database Research Issues in Bioinformatics?" in SSDBM, (Edinburgh, Scotland, UK, 2002), 3.
4. Fields, C., et al. How Many Genes in the Human Genome? Nat Genet, 7 (3). 345-346.
5. <http://www.ncbi.nlm.nih.gov/Genbank/genbankstats.html>. GenBank Statistics, NCBI.
6. Kent, W. J. BLAT: The BLAST-like Alignment Tool. Genome Research, 12 (4). 656-664.
7. McPherson, J. D., et al. A Physical Map of the Human Genome. Nature, 409 (6822). 934-941.
8. Moussouni, F., et al., "Database Challenges for Genome Information in the Post Sequencing Phase", in Proceedings of DEXA'99, (Florence, Italy, 1999), 540-549.
9. Orengo, C. A., et al. From Protein Structure to Function. Curr Opin Struct Biol, 9 (3). 374-382.
10. Purves, W. K. Life, the Science of Biology. Sinauer Associates, Sunderland, MA, 2001.
11. Subramanian, G., et al. Implications of the Human Genome for Understanding Human Biology and Medicine. Jama, 286 (18). 2296-2307.
12. Venter, J. C., et al. Genome Sequence Analysis: Scientific Objectives and Practical Strategies. Trends Biotechnol, 10 (1-2). 8-11.
13. Venter, J. C., et al. The Sequence of the Human Genome. Science, 291 (5507). 1304-1351.