

Editor's Notes

Welcome to the March 2014 issue of the ACM SIGMOD Record!

First, be sure to peruse the Changes to SIGMOD Bylaws, which are notified to ACM members (also through the SIGMOD Record). The changes are due to the fact that the SIGMOD Anthology and the SIGMOD DiSC are no longer produced in their original form, and instead replaced by an online collection. SIGMOD members may object to these changes until July 16, 2014; the details are provided right after these notes.

In the Database Principles column, Andrew McGregor surveys stream algorithms for processing very large graphs. State-of-the art algorithms are surveyed for a variety of popular problems, identifying common algorithmic techniques and the associated smart graph summary data structures, such as spanners and sparsifiers. This very carefully constructed survey can serve as the basis of graduate courses on graph processing algorithms.

The article by Tu, Wang, Zeng and Xu consider the problem of building an energy-efficient, yet performant DBMS. Their approach is based on two strategies: first, energy-aware optimization of queries and storage organization, and second, controlling the power modes of the hardware. The authors developed a system called E²DBMS by augmenting the kernel of PostgreSQL, and report on initial promising results obtained experimenting with their system.

The second article of the issue is by Anglez, Boncz, Larriba-Pey, Fundulaki, Neumann, Erling, Neubauer, Martinez, Kotsev and Toma, which collaborate within the Linked Data Benchmark Council (LDBC) research and development project sponsored by the European Union. The article presents the project objectives and goals, together with an original process and design methodology in designing new benchmarks, with a particular focus on benchmark features challenging the RDF data management engine from a performance perspective.

In the Systems and Prototypes column, Luiz, Ling and Correia describe MITRA, a middleware for replicating DBMSs and making them tolerant to Byzantine faults. Further, MITRA is designed to support transparent management of replicas across DBMSs from several vendors. This feature is interesting when one is to guarantee against the failure of all nodes due to the same bug in the same system. The paper outlines the replication protocol, the system architecture, and presents a brief performance evaluation.

In the Research Centers column, Balazinska, Howe and Suciu present the database research group at the University of Washington. The group topics include scientific data management, big data systems including for instance innovative systems built on top of Dydra and Hadoop, user-facing tools such as the SQLShare project which aims at making the usage of databases simpler by automating or hiding from the users, all aspects but the querying, thus debunking the myth that writing SQL queries is the main issue preventing a wider adoption of SQL. The authors also summarize their research works on probabilistic databases, causality in data transformations and query pricing.

The Open Forum column features two articles. The first, by Pawlish, Varde, Robila and Ranganathan focuses on the energy efficiency of data centers. The authors gathered energy usage data on a real data center belonging to a mid-size university, over three years. A Case-Based Reasoning (CBR) approach is then presented to support deciding which (part of) the operation(s) should be kept in-house and which to delegate to an external cloud, thus forming a hybrid data center.

The second article, by Anciaux, Bouganim, Delot, Illari, Kloul, Mitton and Pucheral is highly original. It considers the problem of providing data services in Least Developed Countries (LDCs), under many

strong restrictions: mostly no internet connection, illiterate users, strong challenges to personal rights and safety through the lack of appropriate law enforcement to which users could appeal in case their personal information is stolen or malevolently used. The authors outline an architecture called Folk-IS, enabling individuals from least developed countries to implement a large distributed system for performing many useful tasks across areas such as banking, culture and leisure services etc.

The issue features two event reports. Wandelt, Deng, Gerdijkov, Mishra, Mitankin, Patil, Siragusa, Tiskin, Wang, Wang and Lese outline an impressive string similarity benchmark effort which they organized under the form of the International Competition on Scalable String Similarity Search and Join (S4) held in conjunction with EDBT/ICDT 2013. The workshop organizers have set up task definitions and a specification of the running environment, then the workshop allowed nine teams to compete on two specific string matching problems, namely k-approximate search and k-approximate join. The paper describes the problems, the methodology, and the teams' results. Second, Darmont and Pedersen report on the Second International Workshop on Cloud Intelligence (Cloud-I 2013), held in conjunction with the VLDB 2013 Conference in Riva del Garda, Italy. Cloud-I featured a keynote on AsterixDB, papers on MapReduce and distributed data provenance, and a panel on challenges for research and industry.

The issue closes with a call for submissions to the First International Workshop on Bringing the Value of "Big Data" to Users (DATA4U), to be held in conjunction with the VLDB conference, in 2014.

Editor-in-Chief change: Yanlei Diao to take over

This issue is a special one for me, since it is the last one I edit! It has been a long journey since the first I had been in charge of, namely March 2010. I am extremely grateful for the honor of serving SIGMOD in this unique position, and I thank first and foremost Yannis Ioannidis, the SIGMOD chair, and SIGMOD Executive committee members of that time, for trusting me with the fate of the Record. I have been fortunate to work with a great team of editors, and to recruit other great ones during these four years. I want to thank and congratulate them again for the patient, careful work in inviting, preparing, overseeing and editing the contributions we are giving back to you, our readers. My thanks go next to the current SIGMOD Executive Committee and in particular the newly elected chairs and co-chairs, Donald Kossmann, Anastasia Ailamaki and Magda Balazinska, who trusted me again with the Record and, when I requested to be relieved of this task, organized the selection process that lead to the designation of the next Editor-in-Chief: Yanlei Diao, from the University of Massachussets.

My best wishes to Yanlei and long live the Record, the most lively database journal I know of (and one with around 70K downloads over the last year: <http://dl.acm.org/citation.cfm?id=J689>)!

Your submissions to the Record are welcome via the submission site:

<http://sigmod.hosting.acm.org/record>

Prior to submitting, be sure to peruse the Editorial Policy on the SIGMOD Record's Web site (<http://www.sigmod.org/publications/sigmod-record/sigmod-record-editorial-policy>).

Ioana Manolescu

March 2014

Past SIGMOD Record Editors:

Harrison R. Morse (1969)
Daniel O'Connell (1971 – 1973)
Randall Rustin (1974-1975)
Douglas S. Kerr (1976-1978)
Thomas J. Cook (1981 – 1983)
Jon D. Clark (1984 – 1985)
Margaret H. Dunham (1986 – 1988)
Arie Segev (1989 – 1995)
Jennifer Widom (1995 – 1996)
Michael Franklin (1996 – 2000)
Ling Liu (2000 – 2004)
Mario Nascimento (2005 – 2007)
Alexandros Labrinidis (2007 – 2009)