THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Report on PODS 2013

Citation for published version:

Fan, W 2013, 'Report on PODS 2013', *ACM SIGACT News*, vol. 44, no. 3, pp. 69-71.
<https://doi.org/10.1145/2527748.2527765>

Digital Object Identifier (DOI):

[10.1145/2527748.2527765](https://doi.org/10.1145/2527748.2527765)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

ACM SIGACT News

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Report on PODS 2013 *

Wenfei Fan
PODS 2013 PC Chair
University of Edinburgh, UK
wenfei@inf.ed.ac.uk

August 1, 2013

The 32nd edition of the ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Databases (PODS) took place from June 24 to June 26, 2013, in New York, USA. The symposium was organized jointly with the ACM SIGMOD International Conference on Management of Data (SIGMOD). PODS focuses on theoretical aspects of data management systems and techniques, and the co-location with SIGMOD stimulates interaction between theory-oriented and system-oriented research.

The proceedings of PODS 2013 is published by ACM Press, and can also be found both on the SIGMOD website (<http://www.sigmod.org>), as well as in the ACM Digital Library (<http://www.acm.org/dl>). The conference program included a keynote talk by Diego Calvanese, two invited tutorials (the first by Piotr Indyk and the second by Pablo Barceló), and 25 contributed papers that were selected by the Program Committee from 97 submissions. Similar to past PODS conferences, most of the contributed papers are preliminary reports on work in progress, and many of them will probably appear in more polished and detailed form in scientific journals.

In his keynote talk, Diego Calvanese highlighted the need for incorporating data modeling features into (business) process modeling languages, and for enriching business process analysis tools to deal with data. He argued that fundamental research has to be conducted on the dichotomy between data and processes. He also identified several key challenges, such as the complexity of verification of data-aware processes when data changes, when the information is incomplete as found in typical knowledge bases, and when the state space is infinite due to the presence of data. The accompanying paper, *Foundations of data-aware process analysis: a database theory perspective* by Diego Calvanese, Giuseppe De Giacomo and Marco Montali, surveys the research on foundations of data-aware (business) processes in more detail, and identifies a number of research directions and open questions.

The tutorial by Piotr Indyk gave an introduction to sketching via hashing, and presented an overview of its applications in data stream algorithms, compressive sensing, numerical

*Editor of the Database Theory Column: Victor Vianu, UC San Diego, vianu@cs.ucsd.edu

linear algebra and sparse Fourier transform. He argued that the method is useful for processing large data sets. The accompanying paper, *Sketching via hashing: from heavy hitters to compressed sensing to sparse Fourier transform* by Piotr Indyk, surveys some of the main applications of sketching via hashing.

The tutorial by Pablo Barceló focused on the expressiveness and complexity of several general-purpose navigational query languages for querying graph databases, namely, regular path queries and its extensions with conjunctions and inverses. He distinguished two semantics for these languages: one based on simple paths, and the other on arbitrary paths. He also presented recent extensions by treating paths as first-class citizens, and by expressing queries that combine the topology of the graph with its underlying data. The accompanying paper *Querying graph databases* by Pablo Barceló gives a comprehensive survey on this subject.

Among the contributed papers, the Program Committee selected the paper *Verification of Database-driven Systems via Amalgamation* by Mikołaj Bojańczyk, Luc Segoufin and Szymon Toruńczyk for the PODS 2013 Best Paper Award. The paper proposes a general framework for inferring decidability and complexity results on the verification problem of database-driven systems. Database-driven systems have been a subject of study in database theory for years, and are of interest to industry, as exemplified by IBM's business artifacts. The paper proves that a general technique based on amalgamation allows to infer upper bounds for the verification problem. It presents several strong results, which generalize and extend previous work on the verification of data-centric systems.

Similarly to preceding years, the contributed papers covered a wide range of subjects, including information retrieval, privacy, meta-querying, data-driven workflows, uncertain data, indexing, query processing and verification, consistent query answering, search languages, RDF and ontology, graph and XML data processing, and dynamic aspects of databases.

The PODS 2013 conference also presented the ACM PODS Alberto O. Mendelzon Test-of-Time Award. The Award is given each year to one or a small number of papers published in the PODS proceedings ten years prior, that had the most impact over the intervening decade. The 2013 Award Committee, constituted of Michael Benedikt, Tova Milo and Dirk Van Gucht, selected the paper *Revealing Information while Preserving Privacy* by Irit Dinur and Kobbi Nissim for this award.

This year the PODS program also included two special events. One was a *Symposium on Theory Challenges in Big Data*, organized by Christopher Ré and Dan Suciu. It took place on the evening of June 23, before the conference. It was a 2-hour event that included 6 invited talks given by Carlos Guestrin, Joe Hellerstein, Andrew McCallum, Muthu Muthukrishnan, Sergei Vassilvitskii and Jeff Ullman, followed by a Q/A session. The talks addressed issues such as whether Big Data is a data management problem? What are the challenges for Data Management introduced by Big Data, and how do they relate to machine learning, statistics, visualization, distributed system? Do these challenges find answers in PODS papers of the past, or do they require new models, techniques and results? This helped to inform the PODS community about a highly relevant emerging topic area. The Symposium was well attended and well received.

The other event was a mini-workshop *Reflections on PODS*, organized by Pablo Barceló and Wim Martens. It took place in the evening of June 25, during the conference. It was a 2.5-hour event that consisted of a panel (80 minutes) and discussions (70 minutes). The panellists were Phokion Kolaitis, Dan Suciu, Tova Milo, Mihalis Yannakakis, Serge Abiteboul, Georg Gottlob, Surajit Chaudhuri and Jeffrey Naughton, who presented their vision on PODS and its future. The talks were very enlightening and entertaining, and the mini-workshop was well attended.

Overall, the program was extremely strong and very interesting. While the number of submissions was low (97, the same as for 1998 and 4 less than 2012), the submissions were unusually strong this year and the conference was very competitive. Many papers with high scores and good reviews could not make it into the final program, due to the scheduling constraints. The co-occurrence with SIGMOD remained beneficial for both conferences. The 2013 SIGMOD/PODS conference was one of the most well-attended in its history, with more than 870 attendees. As in recent years, the significant number of young authors and attendees bodes well for theoretical research in databases.

PODS 2014 will take place in Snowbird, Utah, USA, from June 23 to 25. Richard Hull will be the PODS General Chair and Martin Grohe the PC Chair. More information can be found at <http://www.sigmod2014.org/>.