



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*

## NEWS RELEASE

**Contact:** Jim Ormond  
212-626-0505  
[ormond@hq.acm.org](mailto:ormond@hq.acm.org)

### **2019 ACM SIGMOD/PODS CONFERENCE TO HIGHLIGHT CRITICAL ROLE OF DATA MANAGEMENT IN UNDERPINNING ADVANCED TECHNOLOGIES**

*Conference Will Feature Panels, Keynotes and Tutorials Exploring Vast Importance and Implications of Effective Data Management*

**New York, NY, June 26, 2019** – As advanced technologies continue to proliferate across geographies and industries, data management has become increasingly critical to ensuring that such innovations reach their maximum potential and provide a net benefit to society. In order to address related issues ranging from technical accuracy to practical bias and application, the Association for Computing Machinery Special Interest Group on Management of Data (ACM SIGMOD), together with PODS, the premier international conference on the theoretical aspects of database systems, will host the annual [SIGMOD/PODS 2019](#) conference in Amsterdam from June 30 – July 5, 2019.

The annual conference convenes hundreds of database researchers, practitioners, developers and users to explore cutting-edge ideas and results, and to exchange techniques, tools, and experiences to drive the future of data management forward.

“The breadth and utility of advanced technologies is advancing at breakneck speed today, and many are being employed even as they are being developed” said Anastasia Ailamaki, SIGMOD Program Chair. “Though they may seem disparate – be it virtual reality, artificial intelligence, or others – the vast majority are built upon a foundation of effective data science and data management, so it’s an incredibly exciting time to be a data professional.”

“This year’s PODS program includes a wide selection of research papers which reflect the conference’s tradition of being the premier forum on the theoretical foundations of data management,” added PODS Program Chair Christoph Koch. “At the same time, we continue to broaden our scope. 2019 PODS attendees will learn about the latest research in data management and machine learning, data privacy and security, as well as data ethics.”

### **2019 ACM SIGMOD/PODS HIGHLIGHTS**

#### **Keynote Talks**

**“Differential Privacy and the US Census”**

*Cynthia Dwork, Harvard University*

Differential privacy is a mathematically rigorous definition of privacy tailored to statistical analysis of large datasets. Differentially private systems simultaneously provide useful statistics to the well-intentioned data analyst and strong protection against arbitrarily powerful adversarial system users -- without needing to distinguish between the two. Differentially private systems “don't care” what the adversary knows, now or in the future. Finally, differentially private systems can rigorously bind and control the cumulative privacy loss that accrues over many interactions with the confidential data. These unique properties, together with the abundance of auxiliary data sources and the ease with which they can be deployed by a privacy adversary, led the US Census Bureau to adopt differential privacy as the disclosure avoidance methodology of the 2020 decennial census. This talk will motivate the definition of differential privacy, reflect on the theory-meets-practice experiences of the decennial census, and highlight a few pressing challenges in the field.

### **“Responsible Data Science”**

*Lise Getoor, University of California Santa Cruz*

In this keynote, Getoor will address the challenges data scientists face in creating machine learning and database systems that reason effectively about complex dependencies and uncertainty, and their understanding of the ethical and societal issues involved in data-driven decision-making. She will make a case for the pressing need to integrate algorithmic and statistical principles, social science theories, and basic humanist concepts into data science efforts in order to think critically and constructively about the socio-technical systems being built.

### **“State of Public and Private Blockchains: Myths and Reality”**

*Chandrasekaran Mohan, IBM Almaden Research Center*

This keynote will focus on the technical aspects of blockchains and dispel many myths associated with public blockchains. It has been a decade since the concept of blockchain was invented as the underlying core data structure of the permissionless or public Bitcoin cryptocurrency network. Dr. Mohan considers the adoption and further adaptation of blockchains and smart contracts for use in the permissioned or private environments useful and of practical consequence, and will compare traditional database technologies with blockchain systems' features to identify desirable future research topics.

## **PANEL**

### **SIGMOD Panel on Data Ethics:**

#### **“The Responsibility Challenge for Data”**

*H.V. Jagadish, University of Michigan; Francesco Bonchi, Deputy Director, ISI Foundation; Tina Eliassi-Rad, Northeastern University; Lisa Getoor, University of California, Santa Cruz; Krishna Gummadi, Max Planck Institute for Software Systems; Julia Stoyanovich, New York University*

This panel will discuss the challenges inherent for data scientists in satisfactorily assuming and discharging the broad and critical responsibility of ensuring safe, fair, and private acquisition, presentation, and management of data as increasingly advanced technologies become ubiquitous. While many of the impacts of increasingly ubiquitous data science are beneficial, others may not be, so understanding and managing these impacts is required of every responsible data scientist. Nevertheless, most human decision-makers use algorithms for efficiency purposes and not to make better (i.e., fairer)

decisions; data scientists need to frame the problem with fairness, and other societal impacts, as primary objectives.

## **MODS BEST PAPER AWARDS**

### **SIGMOD Best Paper Award**

#### **“Interventional Fairness: Causal Database Repair for Algorithmic Fairness”**

*Babak Salimi, Luke Rodriguez, Bill Howe, Dan Suciu, University of Washington*

Fairness is increasingly recognized as a critical component of machine learning systems. However, it is the underlying data on which these systems are trained that often reflect discrimination, suggesting a database repair problem. In this paper, the authors propose a method to correctly capture subtle fairness violations.

### **SIGMOD Honorable Mentions**

#### **“Fast general distributed transactions with opacity”**

*Alex Shamis, Microsoft Research; Matthew Renzelmann, Microsoft; Stanko Novakovic, VMware; Georgios Chatzopoulos, EPFL; Aleksandar Dragojevic, Microsoft Research; Dushyanth Narayanan, Microsoft Research; Miguel Castro, Microsoft Research*

Transactions can simplify distributed applications by hiding data distribution, concurrency, and failures from the application developer. This paper extends the design of Fast Remote Memory (FaRM)— which provides strict serializability only for committed transactions — to provide opacity while maintaining FaRM’s high throughput, low latency, and high availability within a modern data center.

#### **“Incremental and Approximate Inference for Faster Occlusion-based Deep CNN Explanations”**

*Supun Nakandala, Arun Kumar and Yannis Papakonstantinou, University of California, San Diego*

Deep Convolutional Neural Networks (CNNs) now match human accuracy in many image prediction tasks, resulting in a growing adoption in e-commerce, radiology, and other domains. In their paper, the authors introduce a novel and comprehensive algebraic framework for incremental CNN inference combining materialized views with multi-query optimization to reduce computational costs.

### **SIGMOD Jim Gray PhD Dissertation Award**

**This award is given to the best database systems dissertation from the previous year.**

#### **“Data Management on Non-Volatile Memory”**

*Joy Arulraj, Assistant Professor of Computing, George Institute of Technology*

We are at an exciting point in the evolution of memory technology. Device manufacturers have created a new non-volatile memory (NVM) technology that can serve as both system memory and storage. The advent of NVM invalidates decades of design decisions that are deeply embedded in today's database management systems (DBMSs). These systems are unable to take full advantage of NVM because their internal architectures are predicated on the assumption that memory is volatile. In his dissertation, Arulraj explores how the core components of a database system can be redesigned for non-volatile memory. He also argues that all types of software systems, including file systems, machine-learning

systems, and key-value stores, are amenable to similar architectural changes to achieve high performance and availability on NVM.

#### **Jim Gray PhD Dissertation Honorable Mention**

##### **“Asynchronous Adventures: Formal Approaches to Querying Big Data in Shared-Nothing Systems”**

*Bas Ketsman, Swiss Federal Institute of Technology, Lausanne (EPFL)*

The central topic addressed by Ketsman’s dissertation, how to efficiently answer queries on large volumes of data in modern parallel and distributed systems, is of major importance in the Big Data era we live in. His thesis provides a clear answer of how forsaking coordination in favor of performance restricts the class of queries that can be evaluated in a distributed manner; as well as how query evaluation can be optimized in the latter coordination-free setting and the setting where one or multiple rounds of synchronization are added. These results not only answer questions that are motivated by real-world problems and systems, but in addition provide clean and simple formalisms that can be followed by theoreticians as well as practitioners. The work therefore represents a major advancement towards providing sorely needed theoretical foundations to the area of query evaluation in modern parallel and distributed systems.

#### **PODS BEST PAPER AWARDS**

##### **PODS Best Student Paper**

##### **“On the Enumeration Complexity of Unions of Conjunctive Queries”**

*Nofar Carmeli, Technion and Markus Kröll, TU Wien*

The authors study the enumeration complexity of Unions of Conjunctive Queries (UCQs). They generalize the notion of free-connexity from Conjunctive Queries (CQs) to UCQs, thus showing that some unions containing intractable CQs are, in fact, tractable.

##### **PODS Best Paper Award**

##### **“Efficient Logspace Classes for Enumeration, Counting, and Uniform Generation”**

*Marcelo Arenas, PUC & IMFD Chile; Luis Alberto Croquevielle, PUC & IMFD Chile; Rajesh Jayaram, Carnegie Mellon University and Cristian Riveros, PUC & IMFD Chile*

In this work, the authors study two simple yet general complexity classes, based on logspace Turing machines, which provide a unifying framework for efficient query evaluation in areas like information extraction and graph databases, among others.

##### **PODS Test of Time Award**

##### **“A General Datalog-Based Framework for Tractable Query Answering over Ontologies”**

*Andrea Cali, Georg Gottlob and Thomas Lukasiewicz, University of Oxford*

Ontologies play a key role in the Semantic Web, data modeling, and information integration. This paper proposes and studies variants of Datalog that are suited for efficient ontological reasoning, and, in particular, for tractable ontology-based query answering.

Additional papers, tutorials, research sessions and in-depth expositions will be presented throughout the multi-day conference. For a complete list of papers and a full schedule of activities, please visit: <https://sigmod2019.org/>.

**About SIGMOD**

[The ACM Special Interest Group on Management of Data \(SIGMOD\)](#) is concerned with the principles, techniques and applications of database management systems and data management technology. Our members include software developers, academic and industrial researchers, practitioners, users, and students. SIGMOD sponsors the annual SIGMOD/PODS conference, one of the most important and selective in the field.

**About ACM**

[ACM, the Association for Computing Machinery](#), is the world's largest educational and scientific computing society, uniting computing educators, researchers and professionals to inspire dialogue, share resources and address the field's challenges. ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

###