

Advances in Database Technology – EDBT 2013

16th International Conference
on Extending Database Technology
Genoa, Italy, March 18-22, 2013
Proceedings

Editors:

Norman W. Paton, University of Manchester, UK

Giovanna Guerrini, Università di Genova, Italy

Barbara Catania, Università di Genova, Italy

Malu Castellanos, HP Labs, USA

Paolo Atzeni, Università Roma Tre, Italy

Piero Fraternali, Politecnico di Milano, Italy

Anastasios Gounaris, Aristotle University of Thessaloniki, Greece

Advances in Database Technology – EDBT 2013
Proceedings of the 16th International Conference
on Extending Database Technology
Genoa, Italy, March 18-22, 2013

Editors:

Norman W. Paton
Giovanna Guerrini
Barbara Catania
Malu Castellanos
Paolo Atzeni
Piero Fraternali
Anastasios Gounaris

The Association for Computing Machinery
2 Penn Plaza, Suite 701
New York, NY, 10121-0701

ACM COPYRIGHT NOTICE. Copyright © 2013 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or permissions@acm.org.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, +1-978-750-8400, +1-978-750-4470 (fax).

ACM ISBN: 978-1-4503-1597-5

Table of Contents

Foreword	ix
Program Committee Members	xi–xiii

Invited Papers

EDBT Keynote 1

Invisible Loading: Access-Driven Data Transfer from Raw Files into Database Systems <i>Azza Abouzied, Daniel J. Abadi, and Avi Silberschatz</i>	1–10
--	------

EDBT Keynote 2

History Repeats Itself: Sensible and NonsenSQL Aspects of the NoSQL Hoopla <i>C. Mohan</i>	11–16
---	-------

Research Sessions

High Performance Query Processing

From A to E: Analyzing TPC’s OLTP Benchmarks - The obsolete, the ubiquitous, the unexplored <i>Pinar Tözün, Ippokratis Pandis, Cansu Kaynak, Djordje Jevdjic, and Anastasia Ailamaki</i>	17–28
---	-------

Query-Aware Compression of Join Results <i>Christopher M. Mullins, Lipyeow Lim, and Christian A. Lang</i>	29–40
--	-------

Web Data Indexing in the Cloud: Efficiency and Cost Reductions <i>Jesús Camacho-Rodríguez, Dario Colazzo, and Ioana Manolescu</i>	41–52
--	-------

Multi-tenant Databases

ProRea - Live Database Migration for Multi-tenant RDBMS with Snapshot Isolation <i>Oliver Schiller, Nazario Cipriani, and Bernhard Mitschang</i>	53–64
---	-------

SWAT: A Lightweight Load Balancing Method for Multitenant Databases <i>Hyun Moon, Hakan Hacigümüş, Yun Chi, and Wang-Pin Hsiung</i>	65–76
--	-------

CloudOptimizer: Multi-tenancy for I/O-Bound OLAP Workloads <i>Hatem A. Mahmoud, Hyun Jin Moon, Yun Chi, Hakan Hacigümüş, Divyakant Agrawal, and Amr El-Abadi</i>	77–88
---	-------

MapReduce

Eagle-Eyed Elephant: Split-Oriented Indexing in Hadoop <i>Mohamed Y. Eltabakh, Fatma Özcan, Yannis Sismanis, Peter J. Haas, Hamid Pirahesh, and Jan Vondrak</i>	89–100
--	--------

Computing n-Gram Statistics in MapReduce <i>Klaus Berberich and Srikanta Bedathur</i>	101–112
--	---------

Processing Multi-way Spatial Joins On Map-Reduce <i>Himanshu Gupta, Bhupesh Chawda, Sumit Negi, Tanveer A. Faruque, L V Subramaniam, and Mukesh Mohania</i>	113–124
Extending Database Technology	
Rapid Experimentation for Testing and Tuning a Production Database Deployment <i>Nedyalko Borisov and Shivnath Babu</i>	125–136
Towards Context-Aware Search and Analysis on Social Media Data <i>Leon R.A. Derczynski, Bin Yang, and Christian S. Jensen</i>	137–142
Proactive Natural Language Search Engine: Tapping into Structured Data on the Web <i>Wensheng Wu</i>	143–148
Anomaly Management using Complex Event Processing <i>Bastian Hoßbach and Bernhard Seeger</i>	149–154
Privacy	
Compromising Privacy in Precise Query Protocols <i>Jonathan L. Dautrich Jr. and Chinya V. Ravishankar</i>	155–166
Efficient Privacy-Aware Record Integration <i>Mehmet Kuzu, Murat Kantarcioglu, Ali Inan, Elisa Bertino, Elizabeth Durham, and Bradley Malin</i>	167–178
Updating Outsourced Anatomized Private Databases <i>Ahmet Erhan Nergiz, Chris Clifton, and Qutaibah M. Malluhi</i>	179–190
Efficient and Accurate Strategies for Differentially-Private Sliding Window Queries <i>Jianneng Cao, Qian Xiao, Gabriel Ghinita, Ninghui Li, Elisa Bertino, and Kian-Lee Tan</i>	191–202
Potpourri	
An Automatic Physical Design Tool for Clustered Column-Stores <i>Alexander Rasin and Stan Zdonik</i>	203–214
Mining Frequent Serial Episodes Over Uncertain Sequence Data <i>Li Wan, Ling Chen, and Chengqi Zhang</i>	215–226
Efficient processing of containment queries on nested sets <i>Ahmed Ibrahim and George H. L. Fletcher</i>	227–238
Inferential Time-Decaying Bloom Filters <i>Jonathan L. Dautrich Jr. and Chinya V. Ravishankar</i>	239–250
Sensors	
Utility-driven Data Acquisition in Participatory Sensing <i>Mehdi Riahi, Thanasis G. Papaioannou, Immanuel Trumme, and Karl Abererr</i>	251–262
An RFID and Particle Filter-Based Indoor Spatial Query Evaluation System <i>Jiao Yu, Wei-Shinn Ku, Min-Te Sun, and Hua Lu</i>	263–274
A Safe Zone Based Approach for Monitoring Moving Skyline Queries <i>Muhammad Aamir Cheema, Xuemin Lin, Wenjie Zhang, and Ying Zhang</i>	275–286

Graph Querying

Compressed Feature-based Filtering and Verification Approach for Subgraph Search <i>Karam Gouda and Mosab Hassaan</i>	287–298
Efficient Query Answering against Dynamic RDF Databases <i>François Goasdoué, Ioana Manolescu, and Alexandra Roatis</i>	299–310
Efficient Breadth-First Search on Large Graphs with Skewed Degree Distributions <i>Haichuan Shang and Masaru Kitsuregawa</i>	311–322

Social Networks and Semantic Querying

CINEMA: Conformity-Aware Greedy Algorithm for Influence Maximization in Online Social Networks <i>Hui Li, Sourav S Bhowmick, and Aixin Sun</i>	323–334
Pollux: Towards Scalable Distributed Real-time Search on Microblogs <i>Liwei Lin, Xiaohui Yu, and Nick Koudas</i>	335–346
Semantic Queries By Example <i>Lipyeow Lim, Haixun Wang, and Min Wang</i>	347–358

Search and Textual Data

Scalable Top-K Spatial Keyword Search <i>Dongxiang Zhang, Kian-Lee Tan, and Anthony K. H. Tung</i>	359–370
Panorama: A Semantic-Aware Application Search Framework <i>Di Jiang, Jan Vosecky, Kenneth Wai-Ting Leung, and Wilfred Ng</i>	371–382
Selectivity Estimation for Hybrid Queries over Text-Rich Data Graphs <i>Andreas Wagner, Veli Bicer, and Thanh D. Tran</i>	383–394

Skyline

Skyline Probability over Uncertain Preferences <i>Qing Zhang, Pengjie Ye, Xuemin Lin, and Ying Zhang</i>	395–405
SkyDiver: A Framework for Skyline Diversification <i>George Valkanas, Apostolos N. Papadopoulos, and Dimitrios Gunopoulos</i>	406–417
Subspace Global Skyline Query Processing <i>Mei Bai, Junchang Xin, and Guoren Wang</i>	418–429

Database as a Service

SWORD: Scalable Workload-Aware Data Placement for Transactional Workloads <i>Abdul Quamar, K. Ashwin Kumar, and Amol Deshpande</i>	430–441
PMAX: Tenant Placement in Multitenant Databases for Profit Maximization <i>Ziyang Liu, Hakan Hacigümüş, Hyun Moon, Yun Chi, and Wang-Pin Hsiung</i>	442–453
Elastic Online Analytical Processing on RAMCloud <i>Christian Tinnefeld, Donald Kossmann, Martin Grund, Joos-Hendrik Boese, Frank Renkes, Vishal Sikka, and Hasso Plattner</i>	454–464

Preference Queries

Skyline Queries in Crowd-Enabled Databases <i>Christoph Lofi, Kinda El Maarry, and Wolf-Tilo Balke</i>	465–476
---	---------

From stars to galaxies: skyline queries on aggregate data <i>Matteo Magnani and Ira Assent</i>	477–488
Efficient Top-k Query Answering using Cached Views <i>Min Xie, Laks V.S. Lakshmanan, and Peter T. Wood</i>	489–500
Stream Query Processing	
Enhanced Stream Processing in a DBMS Kernel <i>Erietta Liarou, Stratos Idreos, Stefan Manegold, and Martin Kersten</i>	501–512
Probabilistic Inference of Object Identifications for Event Stream Analytics <i>Di Wang, Elke Rundensteiner, Richard Ellison, and Han Wang</i>	513–524
High-Performance Complex Event Processing using Continuous Sliding Views <i>Medhabi Ray, Elke A. Rundensteiner, Mo Liu, Chetan Gupta, Song Wang, and Ismail Ari</i>	525–536
Data Integration	
Data Exchange with Arithmetic Operations <i>Balder Ten Cate, Phokion G. Kolaitis, and Walied Othman</i>	537–548
HIL: A High-Level Scripting Language for Entity Integration <i>Mauricio Hernández, Georgia Koutrika, Rajasekar Krishnamurthy, Lucian Popa, and Ryan Wisnesky</i>	549–560
Optimizing Query Rewriting in Ontology-Based Data Access <i>Floriana Di Pinto, Domenico Lembo, Maurizio Lenzerini, Riccardo Mancini, Antonella Poggi, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo</i>	561–572
Industry and Applications Sessions	
Systems and Tools	
Temporal Query Processing in Teradata (short paper) <i>Mohammed Al-Kateb, Ahmad Ghazal, Alain Crolotte, Ramesh Bhashyam, Jaiprakash Chimanchode, and Sai Pavan Pakala</i>	573–578
Near Real-Time Analytics with IBM DB2 Analytics Accelerator <i>Daniel Martin, Oliver Koeth, Johannes Kern, and Iliyana Ivanova</i>	579–588
AppSleuth: a Tool for Database Tuning at the Application Level <i>Wei Cao and Dennis Shasha</i>	589–600
Big Data	
Cost Exploration of Data Sharings in the Cloud <i>Samer Al-Kiswany, Hakan Hacigümüş, Ziyang Liu, and Jagan Sankaranarayanan</i>	601–612
A Performance Comparison of Parallel DBMSs and MapReduce on Large-Scale Text Analytics <i>Fei Chen and Meichun Hsu</i>	613–624
Sparkler: Supporting Large-Scale Matrix Factorization <i>Boduo Li, Sandeep Tata, and Yannis Sismanis</i>	625–636
Applications	
Choosing the Right Crowd: Expert Finding in Social Networks <i>Alessandro Bozzon, Marco Brambilla, Stefano Ceri, Matteo Silvestri, and Giuliano Vesci</i>	637–648

Real-Time Wildfire Monitoring Using Scientific Database and Linked Data Technologies <i>Manolis Koubarakis, Charalambos Kontoes, Stefan Manegold, Manos Karpathiotakis, Kostis Kyzirakos, Konstantina Bereta, George Garbis, Charalampos Nikolaou, Dimitrios Michail, Ioannis Papoutsis, Themistoklis Herekakis, Milena Ivanova, Ying Zhang, Holger Pirk, Martin Kersten, Kallirroi Dogani, Stella Giannakopoulou, and Panayiotis Smeros</i>	649–660
Efficient Multifaceted Screening of Job Applicants <i>Sameep Mehta, Rakesh Pimplikar, Amit Singh, Lav R. Varshney, and Karthik Viswesariah</i>	661–671
Potpourri	
EXLEngine: executable schema mappings for statistical data processing <i>Paolo Atzeni, Bellomarini Luigi, and Bugiotti Francesca</i>	672–682
HyperLogLog in Practice: Algorithmic Engineering of a State of The Art Cardinality Estimation Algorithm <i>Stefan Heule, Marc Nunkesser, and Alexander Hall</i>	683–692
Entity Discovery and Annotation in Tables <i>Gianluca Quercini and Chantal Reynaud</i>	693–704
Demonstration Sessions	
Demonstration Session 1	
iPark: Identifying Parking Spaces from Trajectories <i>Bin Yang, Nicolas Fantini, and Christian S. Jensen</i>	705–708
Limosa: A System for Geographic User Interest Analysis in Twitter <i>Jan Vosecky, Di Jiang, and Wilfred Ng</i>	709–712
Accelerating Spatial Range Queries <i>Alexandros Stougiannis, Thomas Heinis, Farhan Tauheed, and Anastasia Ailamaki</i>	713–716
An Efficient Layout Method for a Large Collection of Geographic Data Entries <i>Sarana Nutanong, Marco D. Adelfio, and Hanan Samet</i>	717–720
In the Mood4: Recommendation by Examples <i>Rubi Boim and Tova Milo</i>	721–724
YmalDB: A Result-Driven Recommendation System for Databases <i>Marina Drosou and Evaggelia Pitoura</i>	725–728
CrowdSeed: Query Processing on Microblogs <i>Zhou Zhao, Wilfred Ng, and Zhijun Zhang</i>	729–732
Hive Open Research Network Platform <i>Jung Hyun Kim, Xilun Chen, K. Selçuk Candan, and Maria Luisa Sapino</i>	733–736
Demonstration Session 2	
Tuning in Action <i>Wei Cao and Dennis Shasha</i>	737–740
PostgreSQL Anomalous Query Detector <i>Bilal Shebaro, Asmaa Sallam, Ashish Kamra, and Elisa Bertino</i>	741–744

Processing XML Queries and Updates on Map/Reduce Clusters <i>Nicole Bidoit, Dario Colazzo, Noor Malla, Maurizio Nolè, Carlo Sartiani, and Federico Ulliana</i>	745–748
CISC: Clustered Image Search by Conceptualization <i>Kaiqi Zhao, Enxun Wei, Qingyu Sui, Kenny Q. Zhu, and Eric Lo</i>	749–752
MinExp-Card: Limiting Data Collection Using a Smart Card <i>Nicolas Anciaux, Walid Bezza, Benjamin Nguyen, and Michalis Vazirgiannis</i>	753–756
PrivComp: A Privacy-aware Data Service Composition System <i>Mahmoud Barhamgi, Djamel Benslimane, Youssef Amghar, Nora Cuppens-Boulahia, and Frederic Cuppens</i>	757–760
ProQua: A System for Evaluating Logic-Based Scoring Functions on Uncertain Relational Data <i>Sebastian Lehrack, Sascha Saretz, and Christian Winkel</i>	761–764
ProvenanceCurious: A Tool to Infer Data Provenance from Scripts <i>Mohammad Rezwanul Huq, Peter M.G. Apers, and Andreas Wombacher</i>	765–768
Tutorials	
Tutorials	
Trust and Reputation in and Across Virtual communities <i>Ehud Gudes and Nurit Gal-Oz</i>	769–772
The W3C PROV family of specifications for modelling provenance metadata <i>Paolo Missier, Khalid Belhajjame, and James Cheney</i>	773–776
Schema Mappings and Data Examples <i>Balder ten Cate, Phokion G. Kolaitis, and Wang-Chiew Tan</i>	777–780

Foreword

Welcome to the 2013 edition of the International Conference on Extending Database Technology (EDBT). This year, EDBT is taking place in Genoa, Italy on March 18-22.

The International Conference on Extending Database Technology is a leading international forum for database researchers, practitioners, developers and users to discuss cutting-edge ideas, and to exchange techniques, tools, and experiences related to data management. Data management provides an essential enabling technology for scientific, engineering, business and social communities. This year's programme provides technical results, experience reports and demonstrations on a variety of important and timely topics, including big data, data integration, privacy and social networks. This year, as in prior years, EDBT is co-located with the International Conference on Database Theory (ICDT).

The programme of EDBT includes keynotes, research papers, industry and application papers, demonstrations and workshops. The proceedings of the workshops are published in a companion volume.

The Research track received 235 submissions, each of which was reviewed by 3 members of the programme committee. A total of 48 papers were accepted, giving an acceptance rate of 20%. The reviewing process included a significant period for discussion of borderline papers, and 5 papers were initially conditionally accepted where the programme committee was keen to see certain presentational refinements in the final versions. One innovation this year was the introduction of 6-Page Extending Database Technology papers, where goal is to enable discussion of visionary ideas, risky approaches or emerging applications that will require extensions to established techniques. There are 3 such papers in the programme, accepted from 22 submissions.

The Industry and Applications track received 24 submissions (more than double the number from the previous year), of which 12 were accepted (11 full and 1 short paper), giving an acceptance rate of 50%. The papers were organized in four sessions: two are mainly industrial papers, one has emphasis on applications and the last one focuses on research impacting concrete practical problems. The reviewing process underwent a heated and rigorous discussion period where every paper that didn't have a unanimous reject was discussed. At the end of this phase a good mix of high quality (according to various criteria) papers was accepted including three conditionally accepted ones that were satisfactorily revised and accepted. The Demonstration track received a record 44 submissions, of which 16 were accepted, giving an acceptance rate of 36%.

The program also includes 3 tutorials on timely topics (trust and reputation in virtual communities, provenance standards, and schema mapping with data examples), and four keynote talks (joint with the ICDT conference) from distinguished researchers from both academia and industry, namely, Jan van den Bussche (Universiteit Hasselt, Belgium), Daniel Abadi (Yale/Hadapt), Luc Segoufin (INRIA & ENS-Cachan) and C. Mohan (IBM).

This year both ICDT and EDBT used EasyChair to support the reviewing process; without such resources, managing the timely reviewing of such large numbers of papers would be very much more laborious. We have also been dependent on many members of the community offering their time in organisational and reviewing roles - we are very grateful for the energy and professionalism they

have exhibited.

We look forward to welcoming you at Palazzo Ducale, historical residence of the Doge in the heart of Genoa, for EDBT 2013. We hope you will enjoy the event and you have a pleasant stay in Genoa.

Norman W. Paton (Program Chair)
Giovanna Guerrini (General Chair)
Barbara Catania (Executive Chair)
Malu Castellanos (Industrial and Applications Chair)
Paolo Atzeni (Tutorial Chair)
Piero Fraternali (Demonstrations Chair)
Anastasios Gounaris (Proceedings Chair)

Program Committee Members

Research

Norman Paton (chair)	Goetz Graefe	Esther Pacitti
Foto Afrati	Torsten Grust	Paolo Papotti
Zohra Bellahsene	Dimitrios Gunopulos	Torben Bach Pedersen
Elisa Bertino	Ralf Hartmut	Jean-Marc Petit
Sourav S Bhowmick	Jiawei Han	Rachel Pottinger
Pedro Bizarro	Melanie Herschel	Sunil Prabhakar
Francesco Bonchi	Arantza Illarramendi	Erhard Rahm
Peter Boncz	Verena Kantere	Louiqa Raschid
Angela Bonifati	Bettina Kemme	Berthold Reinwald
Philippe Bonnet	Alfons Kemper	Mirek Riedewald
Eric Bouillet	Georgia Koutrika	Tore Risch
Alejandro Buchmann	Harumi Kuno	Maria Luisa Sapino
Erik Buchmann	Alberto Laender	Peter Scheuermann
Joos-Hendrik Böse	Georg Lausen	Thomas Seidl
Stefano Ceri	Wang-Chien Lee	Julia Stoyanovich
Lei Chen	Ulf Leser	Letizia Tanca
Yi Chen	Justin Levandoski	Sandeep Tata
Hong Cheng	Xuemin Lin	Evimaria Terzi
Vassilis Christophides	Bin Liu	Jens Teubner
Isabel Cruz	Chengfei Liu	Alejandro Vaisman
Philippe Cudré-Maurou	Bertram Ludaescher	Maurice Van Keulen
x Theodore Dalamagas	Qiong Luo	Maria Esther
Prasad Deshpande	Stefan Manegold	Stratis Viglas
Leonidas Fegaras	Yannis Manolopoulos	Yuqing Wu
Alvaro Fernandes	Marta Mattoso	Jiong Yang
Elena Ferrari	Paolo Missier	Jeffery Xu
Irini Fundulaki	Bernhard Mitschang	Carlo Zaniolo
Johann Gamper	Wilfred Ng	Zhenjie Zhang
Minos Garofalakis	Kjetil Nørkvåg	Shuigeng Zhou
Floris Geerts	Fatma Ozcan	

Industrial and Applications

Malu Castellanos (chair)	Minos Garofalakis	Guy Lohman
Christof Bornhoevd	Zoltan Gyongyi	C Mohan
Edward Chang	Howard Ho	Donovan Schneider
Anish Das Sharma	Meichun Hsu	Haixun Wang
Michael Carey	Stratos Idreos	Hans Zeller
Luna Dong	Vana Kalogeraki	Yu Zheng
Alan Fekete	Christian Konig	

Demonstrations

Piero Fraternali (chair)
Elena Baralis
Andrea Cali
Carlo Curino
Florian Daniel
Petros Daras
Altigran da Silva

Daniel Deutch
Donald Kossman
Geert-Jan Houben
Ioana Manolescu
Davide Martinenghi
Wolfgang Nejdl
Oscar Pastor

Pierre Senellart
Julia Stoyanovich
Fabian M. Suchanek
Ernest Teniente
Riccardo Torlone
Xiaofang Zhou

External Reviewers

Mohammed Abujarour
Reza Akbarinia
Tristan Allard
David Alves
Andrej Andrejev
Fabrizio Angiulli
Daniele Apiletti
Stefan Appel
Roland Assam
Nicola Barbieri
Evandrino Barros
Christian Beecks
Thomas Behr
Jesus Bermudez
George Beskales
Brigitte Boden
Matthias Boehm
Vanessa Braganholo
George Brova
Jianneng Cao
Chen Cao
Tania Cerquitelli
Lijun Chang
Muhammad Cheema
Xu Chu
Eleonora Ciceri
Antonio Corral
Alvaro Coutinho
Victor Cuevas Vicenttin
Chenyun Dai
Victor de Graaff
Gianluca Demartini
Hongbo Deng
Saumen Dey
Jonas Dias
Djellel eddine Difallah
Dora Erdos

Lorena Etcheverry
Alessandro Fiori
Peter M. Fischer
Giorgos Flouris
Maarten Fokkinga
Tobias Freudenreich
Sergej Fries
Sebastian Frischbier
Luca Galli
Junhao Gan
Paolo Garza
Giorgos Giannopoulos
George Giorgidze
Behzad Golshan
Leticia Gomez
Romulo Goncalves
Robert Gottstein
Eric Gribkoff
Luigi Grimaudo
Yann Gripay
Quanquan Gu
Pablo Guerrero
Francesco Gullo
Manish Gupta
Cuneyt Gurcan Akcora
Mohand-Saïd Hacid
Marwan Hassani
Sergio Ilarri
Katerina Ioannou
Anca-Maria Ivanescu
Rohit Jain
Ruoming Jin
Salil Joshi
Wai Kit Wong
Sven Koehler
André Koschmieder
Dimitris Kotsakos

Dimitris Kotzinos
Hardy Kremer
Fabian Laforet
Nikolay Laptev
Jihwan Lee
Max Lehn
Francois Lesueur
Jianxin Li
Rong-Hua Li
Jianxin Li
John Liagouris
Xiang Lian
Bing-Rong Lin
Miguel Liroz-Gistau
Sara Magliacane
Imen Mami
Florent Maseglia
Mirjana Mazuran
Eduardo Mena
Tudor Miu
Mirella M. Moro
Hamid Mousavi
Antoine Ndione
Eduardo Ogasawara
João Oliveirinha
Deepak P
Emanuele Panigati
Apostolos Papadopoulos
Odysseas Papapetrou
Ruggero G. Pensa
Antonio Penta
Eleni Petraki
Holger Pirk
Yoann Pitarch
Vassilis Plachouras
Romila Pradhan
Roman Prokofyev

Bill Province
Martin Przyjaciel-Zablocki
Miao Qiao
Lu Qin
Ba Quan Truong
Gennadi Rabinovitch
Emanuele Rabosio
Angelo Rauso
Robert Rehner
Astrid Rheinlaender
Sean Riddle
Joao B. Rocha-Junior
Natali Ruchansky
Dimitris Sacharidis
Rodrygo Santos
Hassan Sayyadi
Alexander Schätzle
Claudio Schifanella
Oliver Schiller
Tom Schreiber
Fabio Alberto Schreiber
Marian Scuturici
Christian Seeger
Bernhard Seeger

Prithviraj Sen
Marco Serafini
Shuo Shang
Bin Shao
Alexander Shkapsky
Rebecca Simmonds
Dimitrios Skoutas
Tianhong Song
Johannes Starlinger
Alan Stokes
Salmin Sultana
Yizhou Sun
Sabina Surdu
Panagiotis Symeonidis
Roberto Tamassia
Aditya Telang
Christian Thomsen
Eleftherios Tiakas
Konstantin Todorov
Yongxin Tong
Alberto Tonon
Goce Trajcevski
Traian Marius Truta
Kostas Tsihclas

Yuanyuan Tuan
Alexander Ulrich
Patrick Valduries
George Valkanas
Alistair Veitch
Thanasis Vergoulis
Brend Wanders
Wei Wang
Jingjing Wang
Jeroen Weijers
Shanchan Wu
Marcin Wylot
Yin Yang
Xun Yi
Shizhuo Yu
Ye Yuan
Ying Zhang
Wenjie Zhang
Zhiwei Zhang
Xiaofei Zhang
Xiang Zhao
Rui Zhou
Junfeng Zhou
Yuanyuan Zhu