SBBD2016

31º Simpósio Brasileiro de Banco de Dados

04 a 07 de outubro de 2016

SALVADOR – BAHIA – BRASIL

PROCEEDINGS OF THE 31 ST BRAZILIAN SYMPOSIUM ON DATABASES

JAVAM C. MACHADO, CARMEM SATIE HARA, MARTA MATTOSO, VANINHA VIEIRA (ORG.)

SOCIEDADE BRASILEIRA DE COMPUTAÇÃO

Fraunhofer Project Center
for Software and Systems Engineering
at UFBA

In cooperation with

Fraunhofer

UFBA 70 ANOS

computação UFBA

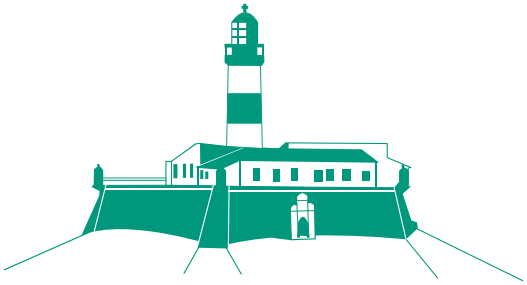

SBBD2016

31th BRAZILIAN SYMPOSIUM ON DATABASES

October 04 to 07, 2016
Salvador – BA – Brazil

PROCEEDINGS

Realization
Brazilian Computer Society – SBC
SBC Special Interest Group on Databases

Organization
Federal University of Bahia (UFBA)
Fraunhofer Project Center in Software and Systems Engineering (FPC@UFBA)

Program Committee Chair
Javam C. Machado (UFC) and Carmem S. Hara (UFPR)

Steering Committee Chair
Vanessa Braganholo (UFF)

Local Organization Chair
Vaninha Vieria (UFBA)

ISSN: 2316-5170

Sistema de Bibliotecas - UFBA

B823p Brazilian Symposium on Databases (31. : 2016 : Salvador, BA)
Proceedings of the 31 st of the Brazilian Symposium on Databases [recurso eletrônico] / Javam C. Machado [et al.] (org.). - Salvador, BA: Brazilian Computer Society, 2016.
287 p.: il.

Disponível em: http://www.sbbd2016.org
ISBN 978-85-7669-342-0

1. Banco de dados - Congressos. 2. Gerência de informações - Congressos. I. Machado, Javam C. II. Título.

CDD - 005.74

Message from the Local Organization Committee Chair

Welcome to the 31st Brazilian Symposium on Databases and to Salvador of Bahia! The Brazilian Symposium on Databases is the official database event of the Brazilian Computer Society (SBC) and the largest venue in Latin America for presentation and discussion of research results in the database domain. Its 31st edition (SBBD 2016) was held in Salvador, the capital of Bahia State, from October 4th to 7th 2016. This year, the symposium was realized in cooperation with ACM-SIGMOD. The local organization was performed by the Federal University of Bahia through the Fraunhofer Project Center for Software and Systems Engineering (FPC-UFBA).

The SBBD 2016 program offers a variety of activities, suited for an audience ranging from undergraduate to Ph.D. students, database professionals, practitioners and researchers. It includes: 3 invited talks and 2 tutorials, presented by distinguished speakers from Brazil and USA; 8 technical sessions, with presentation of high quality researches, selected in a rigorous review process; 4 exciting short courses about hot topics in the area, presented by specialists in their domains; demos and applications session; posters session; the thesis and dissertation workshop, including 11 papers and an invited speaker; and the database systems industry day workshop, with 8 industry cases, an invited speaker and a panel.

SBBD 2016 is the result of the collective effort of a large community, which we gratefully acknowledge. The excellence of its program is a reflex of the competence and effort of the activities chairs. We thank Javam C. Machado (UFC), Carmem S. Hara (UFPR), Marta Mattoso (UFRJ), Eduardo Ogasawara (CEFET-RJ), Sergio Lifschitz (PUC-Rio) and Eduardo C. de Almeida (UFPR). We are extremely grateful to the voluntary, intense, careful and continuous work performed by our colleagues of the Local Organization Committee, in particular to the co-chair João Batista Rocha Jr. (UEFS). We thank the Computer Science Department and the Mathematics and Statistics Institute at UFBA, and our supporting institutions: Universidade Estadual de Feira de Santana (UEFS), Instituto Federal Baiano (IFBaiano), Universidade do Estado da Bahia (UNEB), Universidade Salvador (UNIFACS), SENAI-CIMATEC and Universidade Federal do Rio de Janeiro (UFRJ).

We are also grateful to SBC Board for their trust and organizational support, and to the Steering Committee members for their help, advice and support, particularly to Prof. Vanessa Braganholo (UFF). Further, we thank the program committee members and external reviewers for the high quality reviews, and the authors who submitted their papers to SBBD 2016.

Finally, we would like to express our gratitude to our sponsors: FAPESB, CNPq, CAPES, and the organizations EMC, MTM Tecnologia, PoloIT, Google, Braskem and JusBrasil. Without their support we would not be able to organize this annual

event that brings together our community.

In 30 years, this is the first time that SBBD comes to Bahia. We feel honored with the community trust and we gave our best to organize a great symposium. We wish you all a pleasant week in Salvador, full of rich meetings and discussions. Enjoy the beauty and the flavors of our beloved city and the neighborhood of Rio Vermelho. Axé! Saravá!

Vaninha Vieira, UFBA
SBBD 2016 Local Organization Committee Chair

Table of Contents

31th Brazilian Symposium on Databases

October 04-07, 2016
Salvador – BA – Brazil

SBBD PROCEEDINGS

FULL, SHORT and VISION PAPERS

Promotion

Brazilian Computer Society – SBC
SBC Special Interest Group on Databases

Organization

Fraunhofer Project Center for Software and Systems Engineering – FPC-UFBA
Federal University of Bahia – UFBA

SBBD Program Committee Chair

Javam C. Machado, UFC

SBBD Short Papers Program Co-chair

Carmem S. Hara, UFPR

Editorial

It is a great pleasure to introduce the Proceedings of the Brazilian Symposium on Databases (SBBD) with the full, short and vision papers accepted for presentation at the 31st edition of the symposium, held in Salvador, in the state of Bahia, Brazil, from October 4th to 7th, 2016. This year, SBBD was organized by the Federal University of Bahia (UFBA) in conjunction with the Fraunhofer Project Center for Software and Systems Engineering (FPC-UFBA). SBBD is the official database event of the Brazilian Computer Society (SBC) and the largest venue in Latin America for presentation and discussion of research results in the databases domain. Along with technical sessions, SBBD includes invited talks, tutorials and short courses given by distinguished speakers from the national and international research communities. SBBD regularly promotes a demos and applications session, and a thesis and dissertations workshop as co-located events. This year, SBBD is also promoting a database systems industry day which aims to promote a closer interaction between the industry and the database academic community. For the first time, SBBD had two cycles of submissions for full papers, the first one with deadline by the end of February and the second one by the first week of July. Also this year SBBD had a separate call for short and vision papers with a dedicated committee to review the submissions. All papers submitted to SBBD present interesting results or novel thought-provoking ideas in several subjects on the databases and related areas. For the 2016 edition, SBBD has accepted four kinds of submissions: JIDM articles, full, short and vision papers. Submissions to the JIDM Articles category work in a continuous flow throughout the year. The review is conducted by the editorial board of JIDM, leaded by the editor-in-chief Caetano Traina Jr. Articles accepted by August 8th were invited to be presented at SBBD 2016. This year, only one JIDM articles will be presented in the event.

The review process for the SBBD full papers was performed in one round with a rebuttal phase while for short and vision papers rebuttal was not applied. All authors for full papers were initially notified with the reviews and then had a few days for answering the reviewers' comments during the rebuttal phase. After evaluating the rebuttal comments, a final decision was achieved. Out of 47 submitted papers to both 1st and 2nd cycles, eight papers were accepted as full papers (acceptance rate of 17%), and four as short papers. We also had 55 submitted papers for the short and vision call, 25 of them were accepted (acceptance rate of 45%), including only one as a vision paper. All accepted papers are published in these proceedings and were invited for oral presentation at SBBD 2016. The best full papers will be invited to submit an extended version of the paper to JIDM.

The topics with more accepted submissions (according to the author's selection from the Topics of Interest) were: algorithms and techniques for data mining (ten submissions), concurrency control and recovery (seven submissions), information integration and interoperability (six submissions) and finally access methods, indexing

and new hardware device (five submissions).

The Proceedings of SBBD are the result of the collective effort of a large community, which we gratefully acknowledge. We thank SBBD 2016 local organization committee and its symposium chair, who worked hard to guarantee an outstanding symposium. We are also grateful to the steering committee members for their help, advice and support. We also thank the program committee and short papers committee members, and also external reviewers who provided high quality reviews for the submissions, working very hard in all rounds of review. Finally, we are grateful to the authors who submitted their work to SBBD 2016.

Javam C. Machado (UFC)
SBBD 2016 Program Committee Chair

Carmem S. Hara (UFPR)
SBBD 2016 Short Papers Program Co-chair

31th Brazilian Symposium on Databases

October 04-07, 2016
Salvador – BA – Brazil

Promotion

Brazilian Computer Society – SBC
SBC Special Interest Group on Databases

Organization

Fraunhofer Project Center for Software and Systems Engineering – FPC-UFBA
Federal University of Bahia – UFBA

SBBD Steering Commitee

Altigran Soares da Silva (UFAM)
Caetano Traina Jr. (ICMC - USP)
Carmem S. Hara (UFPR)
Cristina Ciferri (ICMC-USP)
Javam Machado (UFC)
Mirella M. Moro (UFMG)
Vanessa Braganholo (UFF, chair)

SBBD 2016 Commitee

Program Committee Chair
Javam C. Machado (UFC)

Short Papers Program Co-chair
Carmem S. Hara (UFPR)

Local Organization Chairs
Vaninha Vieira (UFBA)

Steering Committee Chair
Vanessa Braganholo (UFF)

Demos and Applications Session Chair
Eduardo Cunha de Almeida (UFPR)

Workshop on Thesis and Dissertations in Databases Chair
Sergio Lifschitz (PUC-Rio)

Lectures Chair
Eduardo Ogasawara (CEFET-RJ)

Tutorials Chair
Marta Mattoso (UFRJ)

Local Organization Committee

Vaninha Vieira (UFBA/FPC-UFBA) (Chair)
Carol Passos (SENAI/FPC-UFBA)
Daniela Claro (UFBA)
Eduardo Jorge (UNEB)
João B. Rocha-Junior (UEFS)
Leila Magalhães (FPC-UFBA)
Marcos Barreto (UFBA)
Paulo Caetano (UNIFACS)
Renato Novais (IFBA/FPC-UFBA)

SBBD Program Committee - Full Papers

Javam C. Machado (UFC, Brazil, PC Chair)
Agma J. Traina (ICMC-USP, Brazil)
Alexandre Plastino (UFF, Brazil)
Alfredo Cuzzocrea (University of Trieste, Italy)
Altigran Silva (UFAM, Brazil)
Amr El Abbadi (UCSB, USA)
Ana Carolina Salgado (UFPE, Brazil)
André C. P. L. F. Carvalho (ICMC-USP, Brazil)
André Santanchè (UNICAMP, Brazil)
Angelo Brayner (UFC, Brazil)
Bernadette Lóscio (UFPE, Brazil)
Bruno Defude (TELECOM SudParis, France)
Caetano Traina Jr. (ICMC-USP, Brazil)
Carina Dorneles (UFSC, Brazil)
Carmem S. Hara (UFPR, Brazil)
Celso Hirata (ITA, Brazil)
Chiara Renso (ISTI-CNR, Italy)
Claudia Roncancio (Grenoble INP, France)
Cláudio de Souza Baptista (UFCG, Brazil)
Clodoveu Davis (UFMG, Brazil)
Cristina Ciferri (ICMC-USP, Brazil)
Daniel Kaster (UEL, Brazil)

Daniel de Oliveira (UFF, Brazil)
Duncan D. Ruiz (PUCRS, Brazil)
Edleno de Moura (UFAM, Brazil)
Eduardo Ogasawara (CEFET-RJ, Brazil)
Elaine P. M. de Sousa (ICMC-USP, Brazil)
Eli Cortez (Microsoft Research, USA)
Fabio Porto (LNCC, Brazil)
Fernanda Baião (UNIRIO, Brazil)
João Eduardo Ferreira (USP, Brazil)
José Antonio Macedo (UFC, Brazil)
Karin Becker (UFRGS, Brazil)
Leo Bertossi (Carleton University, Canada)
Luciano Barbosa (UFPE, Brazil)
Marcos André Gonçalves (UFMG, Brazil)
Maria Claudia Cavalcante (IME, Brazil)
Mario Nascimento (University of Alberta, Canada)
Mirella Moro (UFMG, Brazil)
Mirian Halfeld Ferrari (Université d'Orleans, France)
Oscar Pastor Lopez (Universitat Politecnica de Valencia, Spain)
Pablo Barceló (University of Chile, Chile)
Renata Galante (UFRGS, Brazil)
Ricardo Torres (UNICAMP, Brazil)
Ronaldo dos Santos Mello (UFSC, Brazil)
Sergio Lifschitz (PUC-Rio, Brazil)
Theo Haerder (Un. of Kaiserslautern, Germany)
Valéria Times (UFPE, Brazil)
Vanessa Braganholo (UFF, Brazil)
Victor de Almeida (Petrobrás, Brazil)
Vincent Oria (New Jersey Institute of Technology, USA)
Viviane P. Moreira (UFRGS, Brazil)
Wagner Meira (UFMG, Brazil)

SBBD Program Committee - Short and Vision Papers

Javam C. Machado (UFC, Brazil, PC Chair)
Carmem S. Hara (UFPR, Brazil, Short Papers Co-chair)
Agma J. Traina (ICMC-USP, Brazil)
Alexandre Plastino (UFF, Brazil)
Ana Maria Moura (LNCC, Brazil)
Anderson Ferreira (UFOP, Brazil)
André C. P. L. F. Carvalho (ICMC-USP, Brazil)
André Santanchè (UNICAMP, Brazil)
Angelo Brayner (UFC, Brazil)
Bernadette Lóscio (UFPE, Brazil)
Carina Dorneles (UFSC, Brazil)
Carlos Eduardo Pires (UFCG, Brazil)

Celso Hirata (ITA, Brazil)
Clodoveu Davis (UFMG, Brazil)
Cristiano R. Cervi (UFRGS, Brazil)
Damires Souza (IFPB, Brazil)
Daniel Kaster (UEL, Brazil)
Daniel Notari (UCS, Brazil)
Daniel de Oliveira (UFF, Brazil)
Deise Saccol (UFSM, Brazil)
Denio Duarte (UFFS, Brazil)
Eduardo Cunha de Almeida (UFPR, Brazil)
Eduardo Ogasawara (CEFET-RJ, Brazil)
Elaine P. M. de Sousa (ICMC-USP, Brazil)
Estevam Hruschka Junior (UFSCar, Brazil)
Fabio Porto (LNCC, Brazil)
Fernanda Baião (UNIRIO, Brazil)
Flávio R. C. Sousa (UFC, Brazil)
Helena Graziottin Ribeiro (UCS, Brazil)
João B. Rocha-Junior (UEFS, Brazil)
Jonice Oliveira (UFRJ, Brazil)
José Antonio Macedo (UFC, Brazil)
Jose Maria S. Monteiro Filho (UFC, Brazil)
Karin Becker (UFRGS, Brazil)
Kelly Rosa Braghetto (USP, Brazil)
Luciano Barbosa (UFPE, Brazil)
Luiz Gomes-Junior (UTFPR, Brazil)
Marcos André Gonçalves (UFMG, Brazil)
Maria Camila N. Barioni (UFU, Brazil)
Marilde Santos (UFSCar, Brazil)
Nadia Kozievitch (UTFPR, Brazil)
Raquel Stasiu (PUCPR, Brazil)
Rebeca Schroeder (UDESC, Brazil)
Renata Galante (UFRGS, Brazil)
Renato Bueno (UFSCar, Brazil)
Renato Fileto (UFSC, Brazil)
Ricardo Ciferri (UFSCar, Brazil)
Ricardo Torres (UNICAMP, Brazil)
Robson Cordeiro (ICMC-USP, Brazil)
Robson Fidalgo (UFPE, Brazil)
Ronaldo dos Santos Mello (UFSC, Brazil)
Sergio Lifschitz (PUC-Rio, Brazil)
Sergio Mergen (UFSM, Brazil)
Thiago Silva (UTFPR, Brazil)
Valeria Soares (UFPB, Brazil)
Valéria Times (UFPE, Brazil)
Vania Bogorny (UFSC, Brazil)
Vania Vidal (UFC, Brazil)
Vaninha Vieira (UFBA, Brazil)

Viviane P. Moreira (UFRGS, Brazil)

External Reviewers

Alceu Ferraz Costa (ICMC-USP, Brazil)
André Gonzaga (ICMC-USP, Brazil)
Andre Salvaro Furtado (UFSC, Brazil)
Angelo Frozza (UFSC, Brazil)
Antonio Canabrava Fraideinberze (ICMC-USP, Brazil)
Carlos Andres Ferrero (IFSC, Brazil)
Christian Cesar Bones (ICMC-USP, Brazil)
Davi O. S. de Andrade (UFCG, Brazil)
Eduardo Bezerra (CEFET-RJ, Brazil)
Faisal Nawab (UCSB, USA)
Laercio Gonçalves (CEFET-RJ, Brazil)
Laurence Amaral (UFU, Brazil)
Luciana Romani (ICMC-USP, Brazil)
Maisa Cristina Duarte (UFSCar, Brazil)
Marcelo Iury de Sousa Oliveira (UFRPE, Brazil)
Mauricio Quezada (University of Chile, Chile)
Maxwell Guimarães de Oliveira (UFCG, Brazil)
Raqueline R. M. Penteado (UEM, Brazil)
Saulo D. S. Pedro (UFSCar, Brazil)
Vaibhav Arora (UCSB, USA)

Table of Contents – Full Papers

Table of Contents – Short and Vision Papers

31th Brazilian Symposium on Databases

October 04-07, 2016
Salvador – BA – Brazil

TUTORIALS

Promotion

Brazilian Computer Society – SBC
SBC Special Interest Group on Databases

Organization

Fraunhofer Project Center for Software and Systems Engineering – FPC-UFBA
Federal University of Bahia – UFBA

Tutorials Chair

Marta Mattoso (UFRJ)

Editorial

The tutorials at the Brazilian Symposium on Databases (SBBD) have been a large success over the years. SBBD Tutorials present introductory and advanced discussions on topics within the area of databases. Introductory tutorials target an audience consisting of advanced undergraduate and graduate students, as well as attendees from industry. Hence, they can be thought of as "recycling" courses, though they are not about traditional "textbook" topics. Advanced tutorials, on the other hand, cover a topic's state-of-the-art, motivating and exposing potential research paths.

This year, we have two exciting tutorial themes to be presented at SBBD in Salvador. The first tutorial is entitled "A Practical Tutorial on Deep Learning for Text and Data Integration" and will be presented in Portuguese by Luciano Barbosa, with slide in English. It presents challenges in deep learning for text mining to contribute to the problem of data integration. It covers concepts, the main tools and practical examples with the Keras tool. Luciano Barbosa has recently joined the Computer Science Department at Universidade Federal de Pernambuco (UFPE), Brazil, after an experience as research scientist at IBM Research – Brazil and ATT Research Labs – USA.

The second tutorial, "Provenance in Databases and Scientific Workflows", will be presented by Bertram Ludäscher. The tutorial shows how provenance data describes the lineage and processing history of data as it is transformed through queries or workflows. Provenance data capture, management and analysis is a very active research topic and essential in transparency, reproducibility and preservation of analytical data results. It surveys recent developments in using provenance in databases and scientific workflows motivating for open problems and opportunities for research. Bertram Ludäscher is a professor at the School of Information Sciences at the University of Illinois, Urbana-Champaign, USA, where he directs the Center for Informatics Research in Science and Scholarship, he was previously at the University of California, Davis and the San Diego Supercomputer Center.

I would like to invite you all to attend the tutorials this year in Salvador. Both tutorials cover very hot topics, and we hope you will take advantage of the expertise of the presenters. I'd also like to thank all the authors for contributing to SBBD 2016.

Marta Mattoso (UFRJ)
SBBD 2016 Tutorials Chair

Table of Contents – Tutorials

A Practical Tutorial on Deep Learning for Text and Data Integration

Luciano Barbosa[1]

[1]Centro de Informática
Universidade Federal de Pernambuco (UFPE)
Recife – PE – Brazil

luciano@cin.ufpe.br

Recently, deep learning (DL) techniques have obtained state-of-the-art results in a great variety of text-related problems as, for instance, named entity recognition tagging, sentence classification and machine translation. In this tutorial, we aim to introduce the main DL approaches and their usage for text mining and data integration. For that, we will provide practical examples of deep networks using an easy-to-use Python tool: Keras. At the end of this tutorial, you should be able to build your own deep networks for your specific text-related problems using Keras.

About the Autor

Luciano Barbosa is Assistant Professor in the Computer Science Department at Universidade Federal de Pernambuco. Previously, he worked as Research Scientist in two research labs: IBM Research – Brazil and ATT Research Labs – USA. He obtained his Ph.D. in Computing at University of Utah, and his B.S. and M.S. in Computer Science at Universidade Federal de Pernambuco. His research interests include web mining, text mining and data analytics.

Provenance in Databases and Scientific Workflows

Bertram Ludäscher[1]

[1]Department of Computer Science
University of Illinois
Champaign – IL – USA

ludaesch@illinois.edu

In computer science, data provenance describes the lineage and processing history of data as it is transformed through queries or workflows. Different computer science subdisciplines have studied approaches to capture and exploit provenance, e.g., the systems and programming languages communities. In this tutorial, I will give an overview of basic research questions and results provided by the database and scientific workflow communities. Research in this area ranges from technical studies in database theory (e.g., the use of semi-ring structures to abstract and unify different types of provenance) to more applied techniques (e.g., to efficiently record, store, and query provenance), and various engineering-level questions in-between. Provenance capture and querying capabilities are also playing an increasing role in the reproducibility of scientific workflows, data science applications, the computational sciences. The first half of the tutorial will cover the different uses and types of provenance in scientific workflows, e.g., prospective vs retrospective provenance, and introduce tools that can work with or even combine both forms of provenance to support advanced uses of provenance. In the second half of the tutorial, different notions of provenance in databases will be discussed such as Why-, How-, and Why-Not (missing-answer) provenance. Provenance is a very active research area, and I will end by highlighting some questions and opportunities for future work in databases and workflows.

About the Autor

Bertram Ludäscher is a professor at the School of Information Sciences (iSchool) at the University of Illinois, Urbana-Champaign. At the iSchool, he directs the Center for Informatics Research in Science and Scholarship (CIRSS). He is also a faculty affiliate with the National Center for Supercomputing Applications (NCSA) and the Department of Computer Science. From 2005 to 2014 he was a computer science faculty at the University of California, Davis. His research interests range from scientific data and workflow management, to knowledge representation and reasoning. His current focus includes foundations of provenance and applications, e.g., for automated data quality control and data curation. Until 2004 Ludäscher was a research scientist at the San Diego Supercomputer Center (SDSC) and an adjunct faculty at the CSE Department at UC San Diego. He received his M.S. (Dipl.-Inform.) in computer science from the Technical University of Karlsruhe (K.I.T.) and his PhD from the University of Freiburg, Germany.

31th Brazilian Symposium on Databases

October 04-07, 2016
Salvador – BA – Brazil

INVITED TALKS

Promotion

Brazilian Computer Society – SBC
SBC Special Interest Group on Databases

Organization

Fraunhofer Project Center for Software and Systems Engineering – FPC-UFBA
Federal University of Bahia – UFBA

Table of Contents – Invited Talks

Minimizing Transaction Latency in Geo-replicated Data Stores

Divyakant Agrawal[1]

[1]Department of Computer Science
University of California at Santa Barbara (UCSB)
Santa Barbara – CA – USA

agrawal@cs.ucsb.edu

Web service providers, in general, use NoSQL data stores to ensure scalability and availability of distributed data at the cost of sacrificing transactional guarantees. Recently, major web service providers such as Google have moved towards building storage systems to ensure ACID transactional guarantees for globally distributed data. Spanner from Google uses two-phase locking with Paxos based transaction commitment to ensure consistency of data replicated over multiple datacenters and incurs a large number of cross-datacenter messages. We demonstrate that an alternative approach of replicating commit processing over multiple datacenters can reduce cross-datacenter communication substantially. A common characteristic of both Spanner and the Replicated Commit protocol is that the transaction latencies are subject to round-trip time (RTT) for message communication between the datacenters. We next develop a technique based on log propagation among the datacenters and show that the RTT barrier can be broken subject to certain constraints (e.g., by decoupling consistency and fault-tolerance). A natural question then arises if there is an inherent limit to reducing the commit latency? We answer this question by deriving a lower-bound on commit latency. The sum of the commit latency of any two datacenters is at least the RTT between them. We use the insights and lessons learned while deriving the lower-bound to develop a commit protocol, called Helios, that achieves low commit latencies. Helios actively exchanges transaction logs (history) among the datacenters. The received logs are used to determine whether a transaction can commit or not. The earliest point in the received logs that is needed to commit a transaction is decided by Helios to ensure low commit latency. Also, in a real-world deployment on five datacenters, Helios achieves a commit latency that is close to the optimal.

About the Autor

Divyakant Agrawal is a Professor of Computer Science at the University of California at Santa Barbara. His research expertise is in the areas of database systems, distributed computing, data warehousing, and large-scale information systems. Divy Agrawal has had visiting appointments at IBM Research, NEC Research, ASK.com, Qatar Computing Research Institute, National University of Singapore, and at Google Inc. He has published more than 350 research manuscripts in various forums (journals, conferences, symposia, and workshops) on wide range of topics related to data management and distributed systems and has advised more than 35 Doctoral students during his academic career. He received the 2011 Outstanding Graduate Mentor Award from the Academic Senate at UC Santa Barbara. In 2015, Divy Agrawal and his co-authors received the International Conference on Database Theory Test-of-Time award. His current interests

are in the area of scalable data management and data analysis in Cloud Computing environments, security and privacy of data in the cloud, and scalable analytics over big data. Divy Agrawal is an ACM Distinguished Scientist (2010), an ACM Fellow (2012), and an IEEE Fellow (2012).

The Confounding Problem of Private Data Release

Divesh Srivastava[1]

[1]AT&T Labs-Research
San Ramon – CA – USA

In our Big Data era, as data-driven decision making sweeps through all aspects of society, the demands to make useful data available are growing ever louder. For example, the ubiquity of GPS-enabled devices has resulted in a wealth of data about the movements of individuals and populations, which can be analyzed for useful information to aid in city and traffic planning, disaster preparedness, and so on. But the problem of releasing such data without disclosing confidential information, such as the places people visit, is a subtle and difficult one. Is "private data release" an oxymoron? This talk will delve into the motivations of private data release, explore the challenges, and outline some of the historical and recent approaches developed in response to this confounding problem.

About the Autor

Divesh Srivastava is the head of Database Research at ATT Labs-Research. He is an ACM fellow and the managing editor of the Proceedings of the VLDB Endowment (PVLDB). His research interests and publications span a variety of topics in data management. He received his Ph.D. from the University of Wisconsin, Madison, USA, and his Bachelor of Technology from the Indian Institute of Technology, Bombay, India.

Content-based Similarity Queries on Complex Data – Challenges and Real Applications

Agma J. M. Traina[1]

[1]ICMC – USP
São Carlos – SP – Brazil

agma@icmc.usp.br

The amount and complexity of data generated and managed in nowadays systems, such as images, videos, and time series among others, bring several challenges to the data management developers in order to comply with the expectation of the users and data owners. Not only the majority of the applications demand searching complex data through queries considering several different aspects on the same data, but also getting the answers in a timely manner. Content-based similarity retrieval enables performing queries and analyses using the required features automatically extracted from the data without user intervention.

In this talk we will discuss the challenges posed to the database and related communities in order to provide techniques and tools to overcome the precision and time concerns regarding similarity queries over complex data. Examples and results obtained with two decades long experience over real applications will be presented and discussed.

About the Autor

Agma J. M. Traina is a Professor with the Computer Science Department of the Mathematics and Computer Science Institute at the University of São Paulo at São Carlos. She received her PhD in Computational Physics from the University of São Paulo at São Carlos, Brazil. Agma got her BSc and MSc in Computer Science from the Mathematics and Computer Science Institute at the University of São Paulo at São Carlos. Agma's research interests ranges from complex data indexing and retrieval by content, similarity queries to data visualization and visual data mining. She has focused her research on medical applications supported by image processing techniques, and more recently on climate/agriculture and remote sensing data. Over the years, she has supervised over 40 Graduate students in these areas, and published more than 250 papers in journals and conferences. Agma is a member of the Brazilian Computer Society, ACM and IEEE Computer Society.

Executora

Fraunhofer Project Center
for Software and Systems Engineering
at UFBA

In cooperation with

Fraunhofer

UFBA 70 anos

computação UFBA

Realizado por

SBC

acm In-Cooperation

ACM SIGMOD

Apoio

UEFS

INSTITUTO FEDERAL BAHIA

UNEB

UFRJ

UNIFACS LAUREATE INTERNATIONAL UNIVERSITIES

SENAI CIMATEC SISTEMA FIEB Federação das Indústrias do Estado da Bahia

Fomento

fapesb Fundação de Amparo à Pesquisa do Estado da Bahia

SECRETARIA DE CIÊNCIA, TECNOLOGIA E INOVAÇÃO
BAHIA GOVERNO DO ESTADO

65 anos CAPES

CNPq Conselho Nacional de Desenvolvimento Científico e Tecnológico

Patrocínio Prata:

MTM Tecnologia

Google

POLO IT EXCELÊNCIA EM ALTA DISPONIBILIDADE

Patrocínio Diamante:

EMC²