



# A

## Aberer, Karl

231  Result Selection and Summarization for Web Table Search

1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks

## Agarwal, Manoj K.

1496  Enabling Generic Keyword Search over Raw XML Data

## Aggarwal, Charu C.

185  On Historical Diagnosis of Sensor Streams


1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream


## Agrawal, Rakesh

1597  Big Data: Old Wine in New Bottle?

1600  Data Engineering in Asia: Unique Technical Challenges and Opportunities

## Ailamaki, Anastasia


315  Smooth Scan: Statistics-Oblivious Access Paths

1530  How to Stop Under-Utilization and Love Multicores

## Ait Ouassarah, Azhar

1456  Understanding Business Trends from Data Evolution with Tornado




## Alarabi, Louai

1416  Demonstration of Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs

*A continues on next page...*



## Al-Harathi, Saif

- 1416  Demonstration of Tagheed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs
- 1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data
- 1585  SHAHED: A MapReduce-Based System for Querying and Visualizing Spatio-Temporal Satellite Data


## Ali, Mohamed

- 1215  Predictive Tree: An Efficient Index for Predictive Queries on Road Networks


## Aljubayrin, Saad

- 531  The Safest Path via Safe Zones


## Almeida, Eduardo Cunha de

- 1464  ControVol: A Framework for Controlled Schema Evolution in NoSQL Application Development

## Alturkestani, Tariq

- 843  Automatic Tuning of Bag-of-Tasks Applications



## Aluç, Güneş

- 807  Executing Queries over Schemaless RDF Databases

## Alvarez, Victor

- 1227  A Comparison of Adaptive Radix Trees and Hash Tables


## Alzaidy, Abdulhadi

- 1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data
- 1585  SHAHED: A MapReduce-Based System for Querying and Visualizing Spatio-Temporal Satellite Data

*A continues on next page...*




## An, Aijun

411  Meaningful Keyword Search in Relational Databases with Large and Complex Schema

## Anastasiu, David C.

1549  Understanding Computer Usage Evolution


## Andersen, Ove

1412  EcoSky: Reducing Vehicular Environmental Impact Through Eco-Routing


## Ao, Xiang

891  Online Frequent Episode Mining

## Arasu, Arvind

435  Transaction Processing on Confidential Data Using Cipherbase

## Assent, Ira

1083  Scalable Parallelization of Skyline Computation for Multi-Core Processors

## Averseng, Nicolas

1456  Understanding Business Trends from Data Evolution with Tornado



# B

## Baeuerle, Stefan

1259  Towards a Web-Scale Data Management Ecosystem Demonstrated by SAP HANA


## Bai, Tongxin

1304  Accelerating Big Data Analytics with Collaborative Planning in Teradata Aster 6

## Bakalov, Petko

1364  Time Dependent Transportation Network Models

## Balmin, Andrey

1059  Groupwise Analytics via Adaptive MapReduce

1143  Dynamic Interaction Graphs with Probabilistic Edge Decay

## Bao, Jie


1215  Predictive Tree: An Efficient Index for Predictive Queries on Road Networks

## Bao, Zhifeng


639  Real Time Personalized Search on Social Networks

1504  Elaps: An Efficient Location-Aware Pub/Sub System


## Bao, Zhuowei

375  Answering Regular Path Queries on Workflow Provenance

## Bär, Arian

855  Cache-Oblivious Scheduling of Shared Workloads

## Barber, Ronald

1246  In-Memory BLU Acceleration in IBM's DB2 and dashDB: Optimized for Modern Workloads and Hardware Architectures


*B continues on next page...*




## Barbosa, Denilson

1534  Inferencing in Information Extraction: Techniques and Applications

## Basalamah, Saleh

1416  Demonstration of Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs

1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data


## Berti-Equille, Laure

1440  AllegatorTrack: Combining and Reporting Results of Truth Discovery from Multi-Source Data

## Bertino, Elisa


1011  A Hybrid Private Record Linkage Scheme: Separating Differentially Private Synopses from Matching Records

## Bhangdiya, Amol

1468  The XDa-TA System for Automated Grading of SQL Query Assignments

## Bhowmick, Sourav S.

339  Asymmetric Structure-Preserving Subgraph Queries for Large Graphs

1480  ViSUAL: An HCI-inspired Simulator for Blending Visual Subgraph Query Construction and Processing

1492  PIGEON: Progress Indicator for Subgraph Queries


1500  DAVINCI: Data-Driven Visual Interface Construction for Subgraph Search in Graph Databases

## Bogdanov, Petko


951  Hierarchical In-Network Attribute Compression via Importance Sampling

*B continues on next page...*


**Bøgh, Kenneth S.**

1083  Scalable Parallelization of Skyline Computation for Multi-Core Processors


**Borovica-Gajic, Renata**

315  Smooth Scan: Statistics-Oblivious Access Paths


**Bracher, Shane**

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis

**Bronson, Nathan**

1516  Open Data Challenges at Facebook

**Brunel, Robert**

1280  Supporting Hierarchical Data in SAP HANA

**Buehrer, Gregory**

1071  Towards a Parameter-Free and Parallel Itemset Mining Algorithm in Linearithmic Time

**Burdick, Douglas**

1191  Efficient Sample Generation for Scalable Meta Learning

**Bykau, Siarhei**

1573  Fine-Grained Controversy Detection in Wikipedia



# C

## Cao, Caleb Chen

735  Conflict-Aware Event-Participant Arrangement

## Cao, Jianneng

1011  A Hybrid Private Record Linkage Scheme: Separating Differentially Private Synopses from Matching Records


## Cao, Xin

255  Temporal Spatial-Keyword Top-k Publish/Subscribe


## Cao, Yang

161  Making Pattern Queries Bounded in Big Graphs

## Cao, Yu

1448  PGWinFunc: Optimizing Window Aggregate Functions in PostgreSQL and its Application for Trajectory Data


## Casas, Pedro

855  Cache-Oblivious Scheduling of Shared Workloads

## Cercone, Nick

411  Meaningful Keyword Search in Relational Databases with Large and Complex Schema

## Cerqueus, Thomas

1464  ControVol: A Framework for Controlled Schema Evolution in NoSQL Application Development

## Chakraborty, Tanmoy

555  DiSCern: A Diversified Citation Recommendation System for Scientific Queries


*C continues on next page...*




## Chan, Kam-Lam

- 1376  Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics

## Chandra, Bikash

- 1468  The XDa-TA System for Automated Grading of SQL Query Assignments

## Chang, Lijun

- 387  Diversified Top-K Clique Search

## Chang, Sanghoan

- 1520  The Solid-State Drive Technology, Today and Tomorrow


## Chau, Vinh Tuan

- 1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks

## Chaudhuri, Surajit

- 1  Information at Your Fingertips: Only a Dream for Enterprises?

## Chavan, Shasank

- 1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Chen, Enhong

- 1436  MARS: A Multi-Aspect Recommender System for Point-of-Interest

## Chen, Gang

- 113  PABIRS: A Data Access Middleware for Distributed File Systems

- 591  Efficient Metric Indexing for Similarity Search

- 1268  “Anti-Caching”-based Elastic Memory Management for Big Data

## Chen, Hong






- 1420  CROWN: A Context-Aware Recommender for Web News

*C continues on next page...*







## Chen, Lei

- 6  Cleaning Uncertain Data with a Noisy Crowd
- 279  Answering Why-Not Questions on Spatial Keyword Top-*k* Queries
- 735  Conflict-Aware Event-Participant Arrangement
- 1508  A Graph-Based RDF Triple Store
- 1527  Data-Driven Crowdsourcing: Management, Mining, and Applications


## Chen, Lisi

- 255  Temporal Spatial-Keyword Top-*k* Publish/Subscribe


## Chen, Lu

- 591  Efficient Metric Indexing for Similarity Search
- 1504  Elaps: An Efficient Location-Aware Pub/Sub System

## Chen, Xiao

- 1227  A Comparison of Adaptive Radix Trees and Hash Tables

## Cheng, Jiefeng

- 1131  VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC

## Cheng, Siu-Wing

- 173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees

## Cheng, Xiang

- 1035  Differentially Private Frequent Sequence Mining via Sampling-Based Candidate Pruning

## Chester, Sean

- 1083  Scalable Parallelization of Skyline Computation for Multi-Core Processors

*C continues on next page...*



## Childers, Edwin

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

## Cho, Sangyeun


1520  The Solid-State Drive Technology, Today and Tomorrow

## Choi, Byron


339  Asymmetric Structure-Preserving Subgraph Queries for Large Graphs

1480  ViSUAL: An HCI-inspired Simulator for Blending Visual Subgraph Query Construction and Processing

1492  PIGEON: Progress Indicator for Subgraph Queries

1500  DAVINCI: Data-Driven Visual Interface Construction for Subgraph Search in Graph Databases

## Choi, Dong-Wan

699  Nearest Neighborhood Search in Spatial Databases


## Chon, Jaime

1472  Querying Databases by Snapping Blocks


## Chow, Chi-Yin

1376  Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics

## Christensen, Kenneth Fuglsang

687  Searchlight: Context-Aware Predictive Continuous Querying of Moving Objects in Symbolic Space


## Christiansen, Lasse Linnerup

687  Searchlight: Context-Aware Predictive Continuous Querying of Moving Objects in Symbolic Space


*C continues on next page...*




## Chrysanthis, P.K.

1512  Interactive Preference-Aware Query Optimization


## Chua, Huey Eng

1480  VISUAL: An HCI-inspired Simulator for Blending Visual Subgraph Query Construction and Processing

## Chung, Chin-Wan

699  Nearest Neighborhood Search in Spatial Databases


## Colgan, Maria

1253  Oracle Database In-Memory: A Dual Format In-Memory Database


## Cong, Gao

255  Temporal Spatial-Keyword Top-k Publish/Subscribe

567  A General Graph-Based Model for Recommendation in Event-Based Social Networks

675  SAR: A Sentiment-Aspect-Region Model for User Preference Analysis in Geo-Tagged Reviews


## Cormode, Graham

1023  Conservative or Liberal? Personalized Differential Privacy

## Cui, Bin

267  Finding Top-k Local Users in Geo-Tagged Social Media Data


## Curino, Carlo

1523  Blind Men and an Elephant — Coalescing Open-Source, Academic, and Industrial Perspectives on BigData




# D


## Daghistani, Anas

- 1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data


## Dai, Jian

- 543  Personalized Route Recommendation Using Big Trajectory Data

## Damme, Patrick

- 1460  Enjoy FRDM — Play with a Schema-Flexible RDBMS

## Das, Dinesh

- 1253  Oracle Database In-Memory: A Dual Format In-Memory Database


## Dash, Manoranjan

- 1452  *CDR-To-MoVis*: Developing a *Mobility Visualization* System from CDR Data


## da Silva, Altigran

- 399  Ranking Candidate Networks of Relations to Improve Keyword Search over Relational Databases


## Daudjee, Khuzaima

- 807  Executing Queries over Schemaless RDF Databases

## Davidson, Susan B.

- 375  Answering Regular Path Queries on Workflow Provenance

## De Bra, Paul

- 303  Efficient and Scalable Trie-Based Algorithms for Computing Set Containment Relations

## Decraene, James

- 1452  *CDR-To-MoVis*: Developing a *Mobility Visualization* System from CDR Data


*D continues on next page...*




## Dees, Jonathan

1259  Towards a Web-Scale Data Management Ecosystem Demonstrated by SAP HANA

## De Francisci Morales, Gianmarco

137  The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines


## de Moura, Edleno

399  Ranking Candidate Networks of Relations to Improve Keyword Search over Relational Databases

## Deng, Dong

519  *Two Birds with One Stone: An Efficient Hierarchical Framework for Top-k and Threshold-Based String Similarity Search*


## de Oliveira, Pericles

399  Ranking Candidate Networks of Relations to Improve Keyword Search over Relational Databases

## de Oliveira Jr., Roberto L.

1071  Towards a Parameter-Free and Parallel Itemset Mining Algorithm in Linearithmic Time

## Deshpande, Prasad

999  Indexing and Matching Trajectories Under Inconsistent Sampling Rates

## Deutch, Daniel

1484  se1P: Selective Tracking and Presentation of Data Provenance

## Dey, Akon

125  Scalable Distributed Transactions Across Heterogeneous Stores


## Dhillon, Inderjit S.

1388  Clustering to Forecast Sparse Time-Series Data


*D continues on next page...*



## Ding, Ye

- 1376  Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics

## Ding, Zhiming

- 543  Personalized Route Recommendation Using Big Trajectory Data


## Dittrich, Jens

- 1227  A Comparison of Adaptive Radix Trees and Hash Tables

## Djahandideh, Benjamin

- 1432  CliqueSquare in Action: Flat Plans for Massively Parallel RDF Queries


## Dong, Xin Luna

- 89  Scaling up Copy Detection

## Dong, Zhaoan

- 1408  PandaSearch: A Fine-Grained Academic Search Engine for Research Documents

## Douglas, Chris

- 1523  Blind Men and an Elephant — Coalescing Open-Source, Academic, and Industrial Perspectives on BigData

## Dragut, Eduard C.

- 42  Query-Time Record Linkage and Fusion over Web Databases


## Dvorkin, Igor

- 1167  PerfAugur: Robust Diagnostics for Performance Anomalies in Cloud Services




# E

## Eguro, Ken


435  Transaction Processing on Confidential Data Using Cipherbase


## Eich, Marius

903  Dynamic Programming: The Next Step

## Eldawy, Ahmed

1352  SpatialHadoop: A MapReduce Framework for Spatial Data

1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data

1585  SHAHED: A MapReduce-Based System for Querying and Visualizing Spatio-Temporal Satellite Data


## Eleftheriou, Evangelos

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System


## Elmagarmid, Ahmed K.

42  Query-Time Record Linkage and Fusion over Web Databases

## Ercegovac, Vuk

1059  Groupwise Analytics via Adaptive MapReduce



## Evfimievski, Alexandre

1191  Efficient Sample Generation for Scalable Meta Learning



# F



## Faerber, Franz

- 1259  Towards a Web-Scale Data Management Ecosystem Demonstrated by SAP HANA
- 1280  Supporting Hierarchical Data in SAP HANA

## Faloutsos, Christos

- 1047  HaTen2: Billion-Scale Tensor Decompositions



## Fan, Wei

- 173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees
- 1131  VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC

## Fan, Wenfei

- 161  Making Pattern Queries Bounded in Big Graphs


## Fan, Zhe

- 339  Asymmetric Structure-Preserving Subgraph Queries for Large Graphs
- 1492  PIGEON: Progress Indicator for Subgraph Queries



## Farnan, N.L.

- 1512  Interactive Preference-Aware Query Optimization

## Fekete, Alan

- 125  Scalable Distributed Transactions Across Heterogeneous Stores

## Feng, Jianhua

- 519  *Two Birds with One Stone*: An Efficient Hierarchical Framework for Top-k and Threshold-Based String Similarity Search
- 711  A Location-Aware Publish/Subscribe Framework for Parameterized Spatio-Textual Subscriptions






## Feng, Wei

- 1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream

## Feng, Ziqiang

- 291  Accelerating Aggregation Using Intra-Cycle Parallelism

## Finis, Jan

- 1280  Supporting Hierarchical Data in SAP HANA

## Fischer, Peter M.

- 471  Bi-Temporal Timeline Index: A Data Structure for Processing Queries on Bi-Temporal Data

## Fletcher, George H.L.

- 303  Efficient and Scalable Trie-Based Algorithms for Computing Set Containment Relations


## Foster, Ian

- 1179  LDV: Light-Weight Database Virtualization


## Fournet, Xavier

- 1456  Understanding Business Trends from Data Evolution with Tornado

## Franz, Gerald

- 1280  Supporting Hierarchical Data in SAP HANA

## Franzke, Maximilian

- 1428  A Framework for Computation of Popular Paths from Crowdsourced Data

## Fraser, Campbell

- 315  Smooth Scan: Statistics-Oblivious Access Paths

## Fuhry, David

- 1071  Towards a Parameter-Free and Parallel Itemset Mining Algorithm in Linearithmic Time




# G

## Gal, Avigdor

1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks


## Ganesh, Amit

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Gao, Yunjun

591  Efficient Metric Indexing for Similarity Search

## Garcia-Molina, Hector


5  Data Crowdsourcing: Is it for Real?

195  Comprehensive and Reliable Crowd Assessment Algorithms

219  Entity Resolution with Crowd Errors

627  Identifying Users in Social Networks with Limited Information


## García-Soriano, David


137  The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines

## Ge, Chang

471  Bi-Temporal Timeline Index: A Data Structure for Processing Queries on Bi-Temporal Data

## Ge, Tingjian


747  PIE: Approximate Interleaving Event Matching over Sequences

759  Window-Chained Longest Common Subsequence: Common Event Matching in Sequences




*G continues on next page...*




## Ghanem, Thanaa M.

- 1416  Demonstration of Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs


## Ghani, Sohaib

- 1416  Demonstration of Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs
- 1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data
- 1585  SHAHED: A MapReduce-Based System for Querying and Visualizing Spatio-Temporal Satellite Data

## Ghinita, Gabriel

- 1404  PrivGeoCrowd: A Toolbox for Studying Private Spatial Crowdsourcing


## Gilad, Amir

- 1484  se1P: Selective Tracking and Presentation of Data Provenance


## Gionis, Aristides

- 1155  Bump Hunting in the Dark: Local Discrepancy Maximization on Graphs




## Glavic, Boris

- 1179  LDV: Light-Weight Database Virtualization

## Gleeson, Mike

- 1253  Oracle Database In-Memory: A Dual Format In-Memory Database


## Goasdoué, François

- 771  CliqueSquare: Flat Plans for Massively Parallel RDF Queries
- 1432  CliqueSquare in Action: Flat Plans for Massively Parallel RDF Queries
- 1541  Reasoning on Web Data: Algorithms and Performance

*G continues on next page...*



## Godfrey, Parke

- 411  Meaningful Keyword Search in Relational Databases with Large and Complex Schema


## Goel, Anil K.


- 471  Bi-Temporal Timeline Index: A Data Structure for Processing Queries on Bi-Temporal Data

## Goel, Naman

- 1440  AllegatorTrack: Combining and Reporting Results of Truth Discovery from Multi-Source Data


## Golab, Lukasz

- 855  Cache-Oblivious Scheduling of Shared Workloads

- 879  Size-Constrained Weighted Set Cover

- 1476  SMAS: A Smart Meter Data Analytics System

## Gomes, Joao Bartolo

- 1452  CDR-To-MoVis: Developing a Mobility Visualization System from CDR Data

## Gong, Zhiguo

- 579  Quick-Motif: An Efficient and Scalable Framework for Exact Motif Discovery

## Gottstein, Robert

- 1545  DBMS on Modern Storage Hardware

## Grossniklaus, Michael

- 447  Efficient Structural Bulk Updates on the Pre/Dist/Size XML Encoding

## Grün, Christian

- 447  Efficient Structural Bulk Updates on the Pre/Dist/Size XML Encoding


*G continues on next page...*




## Gu, Ling

- 101  Chronos: An Elastic Parallel Framework for Stream Benchmark Generation and Simulation

## Guan, Xiaohong

- 939  A Tale of Three Graphs: Sampling Design on Hybrid Social-Affiliation Networks

## Guo, Chenjuan

- 543  Personalized Route Recommendation Using Big Trajectory Data

- 1412  EcoSky: Reducing Vehicular Environmental Impact Through Eco-Routing



## Guo, Long

- 1504  Elaps: An Efficient Location-Aware Pub/Sub System



# H

## Haas, Peter J.

- 1059  Groupwise Analytics via Adaptive MapReduce
- 1143  Dynamic Interaction Graphs with Probabilistic Edge Decay

## Haas, Robert

- 1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System


## Hammady, Hossam M.

- 1440  AllegatorTrack: Combining and Reporting Results of Truth Discovery from Multi-Source Data

## Han, Jiawei

- 1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream


## Han, Shuo

- 1508  A Graph-Based RDF Triple Store

## Hanson, Eric N.

- 1239  Evolving the Architecture of SQL Server for Modern Hardware Trends


## Härder, Theo

- 1095  Dynamic Physiological Partitioning on a Shared-Nothing Database Cluster

## Hardock, Sergej

- 1545  DBMS on Modern Storage Hardware


## Hartig, Olaf

- 807  Executing Queries over Schemaless RDF Databases

*H continues on next page...*



## Hase, Sanket


1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Haustein, Nils

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

## He, Cheng

173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees

1131  VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC

## He, Qing

891  Online Frequent Episode Mining

## He, Xiaofeng

483  Dish Comment Summarization Based on Bilateral Topic Analysis

## He, Zhen

531  The Safest Path via Safe Zones

## Hendawi, Abdeltawab M.

1215  Predictive Tree: An Efficient Index for Predictive Queries on Road Networks

## Heng, Wee-Liang

1364  Time Dependent Transportation Network Models

## Hidders, Jan

303  Efficient and Scalable Trie-Based Algorithms for Computing Set Containment Relations


## Hoel, Erik

1364  Time Dependent Transportation Network Models


*H continues on next page...*



## Holloway, Allison

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Holmes, Mark

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis

## Honda, Shinya

1292  AEDSMS: Automotive Embedded Data Stream Management System

## Hou, Jun

615  Robust Clustering of Multi-Type Relational Data via a Heterogeneous Manifold Ensemble

## Hu, Haibo

279  Answering Why-Not Questions on Spatial Keyword Top-*k* Queries

## Hu, Huiqi

711  A Location-Aware Publish/Subscribe Framework for Parameterized Spatio-Textual Subscriptions


## Hu, Songlin

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive

## Hu, Xin

1316  FCCE: Highly Scalable Distributed Feature Collection and Correlation Engine for Low Latency Big Data Analytics

## Hua, Wen

495  Short Text Understanding Through Lexical-Semantic Analysis

## Huai, Jinpeng

161  Making Pattern Queries Bounded in Big Graphs

*H continues on next page...*






## Huang, Feiran

1408  PandaSearch: A Fine-Grained Academic Search Engine for Research Documents

## Huang, Jiamin

963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach

## Huang, Jianbin

1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream


## Huang, Ruizhe

161  Making Pattern Queries Bounded in Big Graphs

## Huang, Shuo

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive


## Huang, Xiaocheng

375  Answering Regular Path Queries on Workflow Provenance



# I


## Idreos, Stratos

315  Smooth Scan: Statistics-Oblivious Access Paths

## Ilyas, Ihab F.

1476  SMAS: A Smart Meter Data Analytics System


## Interlandi, Matteo

18  Proof Positive and Negative in Data Cleaning

## Ishikawa, Yoshiharu

1292  AEDSMS: Automotive Embedded Data Stream Management System

## Islam, Asadul

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis



# J

## Jacobsen, Hans-Arno

- 819  Configurable Hardware-Based Streaming Architecture Using Online Programmable-Blocks
- 1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive





## Jang, Jiyong

- 1316  FCCE: Highly Scalable Distributed Feature Collection and Correlation Engine for Low Latency Big Data Analytics

## Jelitto, Jens

- 1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System


## Jensen, Christian S.

- 279  Answering Why-Not Questions on Spatial Keyword Top-*k* Queries
- 531  The Safest Path via Safe Zones
- 591  Efficient Metric Indexing for Similarity Search
- 1412  EcoSky: Reducing Vehicular Environmental Impact Through Eco-Routing

## Jeon, Inah

- 1047  HaTen2: Billion-Scale Tensor Decompositions

## Jeung, Hoyoung

- 1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis

## Jha, Abhay

- 1388  Clustering to Forecast Sparse Time-Series Data


## Jiang, Jinling

- 267  Finding Top-*k* Local Users in Geo-Tagged Social Media Data

*J continues on next page...*



## Jin, Cheqing

- 1448  PGWinFunc: Optimizing Window Aggregate Functions in PostgreSQL and its Application for Trajectory Data

## Jin, Hai

- 795  Scalable SPARQL Querying Using Path Partitioning

## Jin, Ruoming

- 915  Finding Dense and Connected Subgraphs in Dual Networks



## Jin, Tan

- 927  On Random Walk Based Graph Sampling


## Jo, Insoon

- 1520  The Solid-State Drive Technology, Today and Tomorrow



## Joglekar, Manas

- 195  Comprehensive and Reliable Crowd Assessment Algorithms
- 435  Transaction Processing on Confidential Data Using Cipherbase

## Jorgensen, Zach

- 1023  Conservative or Liberal? Personalized Differential Privacy

## Jossé, Gregor

- 459  Linear Path Skylines in Multicriteria Networks
- 1428  A Framework for Computation of Popular Paths from Crowdsourced Data


## Jurdak, Raja

- 987  Bounded Quadrant System: Error-Bounded Trajectory Compression on the Go




# K

## Kalnis, Panos

843  Automatic Tuning of Bag-of-Tasks Applications

## Kamp, Jesse

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Kang, U


1047  HaTen2: Billion-Scale Tensor Decompositions

## Kantarcioglu, Murat

78  Privacy-Aware Dynamic Feature Selection


1011  A Hybrid Private Record Linkage Scheme: Separating Differentially Private Synopses from Matching Records

## Kaoudi, Zoi

771  CliqueSquare: Flat Plans for Massively Parallel RDF Queries

1432  CliqueSquare in Action: Flat Plans for Massively Parallel RDF Queries


## Kar, Biplab

1468  The XDa-TA System for Automated Grading of SQL Query Assignments

## Kargar, Mehdi

411  Meaningful Keyword Search in Relational Databases with Large and Complex Schema

## Kargın, Yağız

1119  The DBMS — Your Big Data Sommelier

## Karypis, George

1549  Understanding Computer Usage Evolution


*K continues on next page...*




## Kaufmann, Martin

- 471  Bi-Temporal Timeline Index: A Data Structure for Processing Queries on Bi-Temporal Data

## Kaushik, Raghav

- 435  Transaction Processing on Confidential Data Using Cipherbase


## Kawarabayashi, Ken-ichi

- 603  Scalable SimRank Join Algorithm

## Kemper, Alfons

- 1280  Supporting Hierarchical Data in SAP HANA

## Kersten, Martin

- 1119  The DBMS — Your Big Data Sommelier

## Khan, Hina A.

- 327  Progressive Diversification for Column-Based Data Exploration Platforms


## Kingsford, Carl

- 663  Convex Risk Minimization to Infer Networks from Probabilistic Diffusion Data at Multiple Scales

## Kircher, Lukas

- 447  Efficient Structural Bulk Updates on the Pre/Dist/Size XML Encoding

## Koltsidas, Ioannis

- 1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

## Kondo, Derrick

- 1304  Accelerating Big Data Analytics with Collaborative Planning in Teradata Aster 6


## König, Arnd Christian

- 1167  PerfAugur: Robust Diagnostics for Performance Anomalies in Cloud Services


*K continues on next page...*



## Koo, Kee Kiat


1452  *CDR-To-MoVis: Developing a Mobility Visualization System from CDR Data*

## Korn, Flip

879  Size-Constrained Weighted Set Cover


1573  Fine-Grained Controversy Detection in Wikipedia

## Kossmann, Donald


435  Transaction Processing on Confidential Data Using Cipherbase


471  Bi-Temporal Timeline Index: A Data Structure for Processing Queries on Bi-Temporal Data

## Kourtellis, Nicolas

137  The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines

## Koutrika, Georgia

507  Generating Reading Orders over Document Collections


1424  LearningAssistant: A Novel Learning Resource Recommendation System

1538  Goals in Social Media, Information Retrieval and Intelligent Agents


## Krishnaswamy, Shonali Priyadarsini

1452  *CDR-To-MoVis: Developing a Mobility Visualization System from CDR Data*

## Kumar, Manish

1167  PerfAugur: Robust Diagnostics for Performance Anomalies in Cloud Services

## Kusumoto, Mitsuru

603  Scalable SimRank Join Algorithm

## Kusy, Brano

987  Bounded Quadrant System: Error-Bounded Trajectory Compression on the Go


[Search]

[A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) [J](#) [K](#) [L](#) [M](#) [N](#) [O](#) [P](#) [Q](#) [R](#) [S](#) [T](#) [U](#) [V](#) [W](#) [X](#) [Y](#) [Z](#)



# L


Lahiri, Tirthankar

1253  Oracle Database In-Memory: A Dual Format In-Memory Database


Larson, Per-Ake

1239  Evolving the Architecture of SQL Server for Modern Hardware Trends


Lee, A.J.

1512  Interactive Preference-Aware Query Optimization

Lee, Dongwon

1527  Data-Driven Crowdsourcing: Management, Mining, and Applications

Lee, Teck-Hua

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

Lehner, Wolfgang

1259  Towards a Web-Scale Data Management Ecosystem Demonstrated by SAP HANA

1460  Enjoy FRDM — Play with a Schema-Flexible RDBMS

Lento, Thomas

1516  Open Data Challenges at Facebook


Li, Chengkai

891  Online Frequent Episode Mining

Li, Cuiping

1420  CROWN: A Context-Aware Recommender for Web News

Li, Feng





879  Size-Constrained Weighted Set Cover

*L continues on next page...*





## Li, Guoliang

- 519  *Two Birds with One Stone: An Efficient Hierarchical Framework for Top-k and Threshold-Based String Similarity Search*
- 639  Real Time Personalized Search on Social Networks
- 711  A Location-Aware Publish/Subscribe Framework for Parameterized Spatio-Textual Subscriptions
- 1504  Elaps: An Efficient Location-Aware Pub/Sub System

## Li, Jia

- 1408  PandaSearch: A Fine-Grained Academic Search Engine for Research Documents

## Li, Lin

- 1436  MARS: A Multi-Aspect Recommender System for Point-of-Interest


## Li, Lu

- 66  Efficient Secure Similarity Computation on Encrypted Trajectory Data

## Li, Rong-Hua

- 927  On Random Walk Based Graph Sampling

## Li, Xian

- 89  Scaling up Copy Detection

## Li, Xiaoli

- 831  Network Motif Discovery: A GPU Approach
- 1452  *CDR-To-MoVis: Developing a Mobility Visualization System from CDR Data*

## Li, Xin

- 1436  MARS: A Multi-Aspect Recommender System for Point-of-Interest

## Li, Xinhao

- 591  Efficient Metric Indexing for Similarity Search

*L continues on next page...*



Li, Xutao

567  A General Graph-Based Model for Recommendation in Event-Based Social Networks


Li, Yanhua

1376  Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics

Li, Youhuan

1508  A Graph-Based RDF Triple Store


Li, Yuchen

639  Real Time Personalized Search on Social Networks

Li, Yuhong


579  Quick-Motif: An Efficient and Scalable Framework for Exact Motif Discovery

Li, Zheng


747  PIE: Approximate Interleaving Event Matching over Sequences

Li, Zhenguo

173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees

1131  VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC


Li, Zhengyi

1035  Differentially Private Frequent Sequence Mining via Sampling-Based Candidate Pruning

Liang, Ying

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive

Liarou, Erietta

1530  How to Stop Under-Utilization and Love Multicores

*L continues on next page...*




## Lightstone, Sam

- 1246  In-Memory BLU Acceleration in IBM's DB2 and dashDB: Optimized for Modern Workloads and Hardware Architectures

## Lin, Wenqing


- 831  Network Motif Discovery: A GPU Approach


## Lin, Xin

- 279  Answering Why-Not Questions on Spatial Keyword Top-k Queries

## Lin, Xuemin

- 30  Cleaning Structured Event Logs: A Graph Repair Approach

- 387  Diversified Top-K Clique Search

- 1107  AP-Tree: Efficiently Support Continuous Spatial-Keyword Queries over Stream

## Ling, Tok Wang


- 1408  PandaSearch: A Fine-Grained Academic Search Engine for Research Documents

## Liu, An

- 66  Efficient Secure Similarity Computation on Encrypted Trajectory Data


- 351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs

## Liu, Guanfeng

- 66  Efficient Secure Similarity Computation on Encrypted Trajectory Data

- 351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs

## Liu, Jiajun

- 423  Interactive Top-k Spatial Keyword Queries


- 987  Bounded Quadrant System: Error-Bounded Trajectory Compression on the Go

*L continues on next page...*



## Liu, Lei

507  Generating Reading Orders over Document Collections

1424  LearningAssistant: A Novel Learning Resource Recommendation System


## Liu, Ling

795  Scalable SPARQL Querying Using Path Partitioning

## Liu, Qiang

1420  CROWN: A Context-Aware Recommender for Web News

## Liu, Qin

1131  VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC

## Liu, Wantao

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive

## Liu, Xiufeng

1476  SMAS: A Smart Meter Data Analytics System

## Liu, Yiqun

711  A Location-Aware Publish/Subscribe Framework for Parameterized Spatio-Textual Subscriptions


## Liu, Zheng

6  Cleaning Uncertain Data with a Noisy Crowd

## Lo, Eric

291  Accelerating Aggregation Using Intra-Cycle Parallelism

## Loaiza, Juan

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

*L continues on next page...*




## Lohman, Guy

- 1246  In-Memory BLU Acceleration in IBM's DB2 and dashDB: Optimized for Modern Workloads and Hardware Architectures

## Lomet, David B.

- 1203  High Performance Temporal Indexing on Modern Hardware


## Lu, Hua


- 267  Finding Top-k Local Users in Geo-Tagged Social Media Data

## Lu, Jiaheng


- 1408  PandaSearch: A Fine-Grained Academic Search Engine for Research Documents

## Lui, John C.S.

- 939  A Tale of Three Graphs: Sampling Design on Hybrid Social-Affiliation Networks

- 1131  VENUS: Vertex-Centric Streamlined Graph Computation on a Single PC

## Luo, Ge

- 173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees

## Luo, Jun

- 1376  Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics


## Luo, Ping

- 891  Online Frequent Episode Mining

## Luo, Yongming

- 303  Efficient and Scalable Trie-Based Algorithms for Computing Set Containment Relations


## Lyons, Kenneth B.

- 89  Scaling up Copy Detection



# M


## Ma, Jiansong

- 1448  PGWinFunc: Optimizing Window Aggregate Functions in PostgreSQL and its Application for Trajectory Data


## Macko, Peter

- 363  LLAMA: Efficient Graph Analytics Using Large Multiversioned Arrays


## Macnaughton, Neil

- 1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Maehara, Takanori

- 603  Scalable SimRank Join Algorithm

## Magdy, Amr

- 1416  Demonstration of Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs


## Malik, Tanu

- 1179  LDV: Light-Weight Database Virtualization

## Malin, Bradley

- 78  Privacy-Aware Dynamic Feature Selection

## Manegold, Stefan

- 1119  The DBMS — Your Big Data Sommelier




## Mani, Murali

- 783  INSURE: An Integrated Load Reduction Framework for XML Stream Processing


*M continues on next page...*



## Manolescu, Ioana

- 771  CliqueSquare: Flat Plans for Massively Parallel RDF Queries
- 1432  CliqueSquare in Action: Flat Plans for Massively Parallel RDF Queries
- 1541  Reasoning on Web Data: Algorithms and Performance

## Mansour, Essam

- 843  Automatic Tuning of Bag-of-Tasks Applications

## Mao, Rui

- 927  On Random Walk Based Graph Sampling

## Marathe, Virendra J.

- 363  LLAMA: Efficient Graph Analytics Using Large Multiversed Arrays


## Margo, Daniel W.

- 363  LLAMA: Efficient Graph Analytics Using Large Multiversed Arrays

## Markl, Volker

- 1191  Efficient Sample Generation for Scalable Meta Learning



## Marwah, Vineet

- 1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Mathioudakis, Michael

- 1155  Bump Hunting in the Dark: Local Discrepancy Maximization on Graphs


## May, Norman

- 471  Bi-Temporal Timeline Index: A Data Structure for Processing Queries on Bi-Temporal Data
- 1280  Supporting Hierarchical Data in SAP HANA


*M continues on next page...*




## McClenaghan, Michael

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis


## Meng, Weiyi


89  Scaling up Copy Detection

## Miklós, Zoltán


1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks

## Milo, Tova

375  Answering Regular Path Queries on Workflow Provenance

1527  Data-Driven Crowdsourcing: Management, Mining, and Applications


## Mischewski, Liam

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis

## Modani, Natwar

555  DiSCern: A Diversified Citation Recommendation System for Scientific Queries


## Moerkotte, Guido


903  Dynamic Programming: The Next Step


## Mokbel, Mohamed F.

1215  Predictive Tree: An Efficient Index for Predictive Queries on Road Networks

1352  SpatialHadoop: A MapReduce Framework for Spatial Data

1416  Demonstration of Taghreed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs


1444  A Demonstration of Shahed: A MapReduce-Based System for Querying and Visualizing Satellite Data

1585  SHAHED: A MapReduce-Based System for Querying and Visualizing Spatio-Temporal Satellite Data






## Moskovitch, Yuval

1484  selP: Selective Tracking and Presentation of Data Provenance


## Mu, Yadong

173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees


## Mukherjee, Arko Provo

243  Mining Maximal Cliques from an Uncertain Graph


## Mukherjee, Niloy

1253  Oracle Database In-Memory: A Dual Format In-Memory Database


## Mullick, Atrayee

1253  Oracle Database In-Memory: A Dual Format In-Memory Database


## Musleh, Mashaal

1416  Demonstration of Tagheed: A System for Querying, Analyzing, and Visualizing Geotagged Microblogs

## Muthulingam, Sujatha

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Mylopoulos, John

1538  Goals in Social Media, Information Retrieval and Intelligent Agents




# N

## Nagar, Seema

555  DiSCern: A Diversified Citation Recommendation System for Scientific Queries

## Najafi, Mohammadreza

819  Configurable Hardware-Based Streaming Architecture Using Online Programmable-Blocks


## Nakamoto, Yukikazu

1292  AEDSMS: Automotive Embedded Data Stream Management System


## Narayanam, Ramasuri

555  DiSCern: A Diversified Citation Recommendation System for Scientific Queries

## Nascimento, Mario A.

1428  A Framework for Computation of Popular Paths from Crowdsourced Data


## Nasir, Muhammad Anis Uddin

137  The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines


## Nawab, Faisal

1203  High Performance Temporal Indexing on Modern Hardware


## Nayak, Richi

615  Robust Clustering of Multi-Type Relational Data via a Heterogeneous Manifold Ensemble

## Neuber, Glenn

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis


## Neumann, Thomas

1280  Supporting Hierarchical Data in SAP HANA



*N continues on next page...*





## Nguyen, Hong H.

- 1500  DAVINCI: Data-Driven Visual Interface Construction for Subgraph Search in Graph Databases

## Nguyen, Quoc Viet Hung

- 231  Result Selection and Summarization for Web Table Search
- 1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks

## Nguyen, Thanh Tam

- 231  Result Selection and Summarization for Web Table Search
- 1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks


## Norwood, Anjali

- 1304  Accelerating Big Data Analytics with Collaborative Planning in Teradata Aster 6




# O

Ong, N.R.

1512  Interactive Preference-Aware Query Optimization

Ooi, Beng Chin

1268  “Anti-Caching”-based Elastic Memory Management for Big Data


Orgun, Mehmet A.

351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs

Ouzzani, Mourad

42  Query-Time Record Linkage and Fusion over Web Databases

Özsu, M. Tamer

807  Executing Queries over Schemaless RDF Databases

1508  A Graph-Based RDF Triple Store



# P

P., Deepak

999  Indexing and Matching Trajectories Under Inconsistent Sampling Rates

Pandit, Aditi

1304  Accelerating Big Data Analytics with Collaborative Planning in Teradata Aster 6

Papadimitrou, Dimitra

1538  Goals in Social Media, Information Retrieval and Intelligent Agents

Papalexakis, Evangelos E.

1047  HaTen2: Billion-Scale Tensor Decompositions

Parameswaran, Aditya

195  Comprehensive and Reliable Crowd Assessment Algorithms

Parthasarathy, Srinivasan

1071  Towards a Parameter-Free and Parallel Itemset Mining Algorithm in Linearithmic Time


Pattuk, Erman

78  Privacy-Aware Dynamic Feature Selection

Pease, David

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

Pedersen, Torben Bach

687  Searchlight: Context-Aware Predictive Continuous Querying of Moving Objects in Symbolic Space

Pei, Jian

30  Cleaning Structured Event Logs: A Graph Repair Approach


*P continues on next page...*



## Pei, Xubin

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive

## Peng, Liping

1059  Groupwise Analytics via Adaptive MapReduce

## Petermann, Martin

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

## Petit, Jean-Marc

1456  Understanding Business Trends from Data Evolution with Tornado

## Petrov, Ilia

1545  DBMS on Modern Storage Hardware


## Pham, Quan

1179  LDV: Light-Weight Database Virtualization


## Pham, Tuan-Anh Nguyen

567  A General Graph-Based Model for Recommendation in Event-Based Social Networks


## Pihl, Jeppe

687  Searchlight: Context-Aware Predictive Continuous Querying of Moving Objects in Symbolic Space


## Pirk, Holger

1119  The DBMS — Your Big Data Sommelier

## Porobic, Danica

1530  How to Stop Under-Utilization and Love Multicores

## Psaroudakis, Iraklis

1530  How to Stop Under-Utilization and Love Multicores




# Q

## Qi, Jianzhong


531  The Safest Path via Safe Zones

## Qin, Lu

387  Diversified Top-K Clique Search

927  On Random Walk Based Graph Sampling

## Quiané-Ruiz, Jorge-Arnulfo

771  CliqueSquare: Flat Plans for Massively Parallel RDF Queries

1432  CliqueSquare in Action: Flat Plans for Massively Parallel RDF Queries




# R


Rabl, Tilmann

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive


Radhakrishnan, Bharath

1468  The XDa-TA System for Automated Grading of SQL Query Assignments

Raghavan, Sriram

999  Indexing and Matching Trajectories Under Inconsistent Sampling Rates


Raja, Vivekanandhan

1253  Oracle Database In-Memory: A Dual Format In-Memory Database


Ramamritham, Krithi

1496  Enabling Generic Keyword Search over Raw XML Data

Ramamurthy, Ravi

435  Transaction Processing on Confidential Data Using Cipherbase


Raman, Vijayshankar

1246  In-Memory BLU Acceleration in IBM's DB2 and dashDB: Optimized for Modern Workloads and Hardware Architectures

Ranu, Sayan

999  Indexing and Matching Trajectories Under Inconsistent Sampling Rates

Rao, Fang-Yu







1011  A Hybrid Private Record Linkage Scheme: Separating Differentially Private Synopses from Matching Records

Rashid, Al M.


1549  Understanding Computer Usage Evolution

*R continues on next page...*




**Ray, Shubhankar**1388  Clustering to Forecast Sparse Time-Series Data**Reddy, K.V. Maheshwara**1468  The XDa-TA System for Automated Grading of SQL Query Assignments**Reinwald, Berthold**1191  Efficient Sample Generation for Scalable Meta Learning**Renz, Matthias**1428  A Framework for Computation of Popular Paths from Crowdsourced Data**Revol, Romain**1456  Understanding Business Trends from Data Evolution with Tornado**Rezig, El Kindi**42  Query-Time Record Linkage and Fusion over Web Databases**Richter, Stefan**1227  A Comparison of Adaptive Radix Trees and Hash Tables**Ricketts, Daniel**1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis**Roatiş, Alexandra**1541  Reasoning on Web Data: Algorithms and Performance**Röhm, Uwe**125  Scalable Distributed Transactions Across Heterogeneous Stores**Rojcewicz, S.E.**1512  Interactive Preference-Aware Query Optimization*R continues on next page...*


**Roth, Marty**

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

**Roy, Sudip**

1167  PerfAugur: Robust Diagnostics for Performance Anomalies in Cloud Services

**Ruehrup, Stefan**

855  Cache-Oblivious Scheduling of Shared Workloads


**Rundensteiner, Elke A.**


783  INSURE: An Integrated Load Reduction Framework for XML Stream Processing




# S

## Sadiq, Shazia


963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach

975  Approximate Keyword Search in Semantic Trajectory Database


## Sadoghi, Mohammad

819  Configurable Hardware-Based Streaming Architecture Using Online Programmable-Blocks

## Saha, Barna

879  Size-Constrained Weighted Set Cover

## Sahli, Majed

843  Automatic Tuning of Bag-of-Tasks Applications

## Sailer, Reiner

1316  FCCE: Highly Scalable Distributed Feature Collection and Correlation Engine for Low Latency Big Data Analytics


## Sarafijanovic, Slavisa

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

## Sato, Kenya

1292  AEDSMS: Automotive Embedded Data Stream Management System

## Schales, Douglas L.

1316  FCCE: Highly Scalable Distributed Feature Collection and Correlation Engine for Low Latency Big Data Analytics


## Schall, Daniel

1095  Dynamic Physiological Partitioning on a Shared-Nothing Database Cluster


*S continues on next page...*




## Schelter, Sebastian

1191  Efficient Sample Generation for Scalable Meta Learning


## Scherzinger, Stefanie

1464  ControVol: A Framework for Controlled Schema Evolution in NoSQL Application Development

## Schiavone, Mirko

855  Cache-Oblivious Scheduling of Shared Workloads

## Schiefer, Berni

1246  In-Memory BLU Acceleration in IBM's DB2 and dashDB: Optimized for Modern Workloads and Hardware Architectures

## Scholl, Marc H.

447  Efficient Structural Bulk Updates on the Pre/Dist/Size XML Encoding

## Schubert, Matthias

459  Linear Path Skylines in Multicriteria Networks

## Scuturici, Vasile-Marian

1456  Understanding Business Trends from Data Evolution with Tornado

## Seaman, Brian

1388  Clustering to Forecast Sparse Time-Series Data

## Sefer, Emre

663  Convex Risk Minimization to Infer Networks from Probabilistic Diffusion Data at Multiple Scales

## Seipp, Harald

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System


*S continues on next page...*




## Seltzer, Margo I.

363  LLAMA: Efficient Graph Analytics Using Large Multiversed Arrays


## Serafini, Marco

137  The Power of Both Choices: Practical Load Balancing for Distributed Stream Processing Engines


## Shah, Shetal

1468  The XDa-TA System for Automated Grading of SQL Query Assignments


## Shahabi, Cyrus

1404  PrivGeoCrowd: A Toolbox for Studying Private Spatial Crowdsourcing

## Shan, Ming-Chien


101  Chronos: An Elastic Parallel Framework for Stream Benchmark Generation and Simulation

## Shang, Shuo

423  Interactive Top- $k$  Spatial Keyword Queries

987  Bounded Quadrant System: Error-Bounded Trajectory Compression on the Go

## Sharaf, Mohamed A.

327  Progressive Diversification for Column-Based Data Exploration Platforms

## She, Jieying

735  Conflict-Aware Event-Participant Arrangement

## Shekelyan, Michael

459  Linear Path Skylines in Multicriteria Networks


## Shen, Xuchuan

1508  A Graph-Based RDF Triple Store

*S continues on next page...*



## Shi-Nash, Amy

1452  *CDR-To-MoVis: Developing a Mobility Visualization System from CDR Data*


## Shkapsky, Alexander

867  *Optimizing Recursive Queries with Monotonic Aggregates in DeALS*

## Shun, Julian

149  *Multicore Triangle Computations Without Tuning*


## Šidlauskas, Darius

1083  *Scalable Parallelization of Skyline Computation for Multi-Core Processors*

## Sidle, Richard

1246  *In-Memory BLU Acceleration in IBM's DB2 and dashDB: Optimized for Modern Workloads and Hardware Architectures*

## Silva, Arlei

951  *Hierarchical In-Network Attribute Compression via Importance Sampling*

## Silva, Yasin N.

1472  *Querying Databases by Snapping Blocks*

## Simmen, David

1304  *Accelerating Big Data Analytics with Collaborative Planning in Teradata Aster 6*

## Simske, Steve

507  *Generating Reading Orders over Document Collections*


## Singh, Ambuj K.

951  *Hierarchical In-Network Attribute Compression via Importance Sampling*

*S continues on next page...*




## Sismanis, Yannis

1059  Groupwise Analytics via Adaptive MapReduce

1143  Dynamic Interaction Graphs with Probabilistic Edge Decay


## Skoumas, Georgios

1428  A Framework for Computation of Popular Paths from Crowdsourced Data

## Sommer, Philipp

987  Bounded Quadrant System: Error-Bounded Trajectory Compression on the Go

## Song, Chunyao

759  Window-Chained Longest Common Subsequence: Common Event Matching in Sequences


## Song, Shaoxu

30  Cleaning Structured Event Logs: A Graph Repair Approach


## Soto, Juan


1191  Efficient Sample Generation for Scalable Meta Learning

## Soylomez, Ekrem

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Srivastava, Divesh

89  Scaling up Copy Detection

879  Size-Constrained Weighted Set Cover

1573  Fine-Grained Controversy Detection in Wikipedia


## Stoecklin, Marc Ph.


1316  FCCE: Highly Scalable Distributed Feature Collection and Correlation Engine for Low Latency Big Data Analytics

*S continues on next page...*




## Su, Han


423  Interactive Top-*k* Spatial Keyword Queries

963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach


## Su, Sen

1035  Differentially Private Frequent Sequence Mining via Sampling-Based Candidate Pruning

## Sudarshan, S.

1468  The XDa-TA System for Automated Grading of SQL Query Assignments

## Szlichta, Jaroslaw

411  Meaningful Keyword Search in Relational Databases with Large and Complex Schema





# T

## Tagarelli, Andrea

1549  Understanding Computer Usage Evolution

## Takada, Hiroaki

1292  AEDSMS: Automotive Embedded Data Stream Management System

## Tan, Kian-Lee

255  Temporal Spatial-Keyword Top-k Publish/Subscribe

639  Real Time Personalized Search on Social Networks

711  A Location-Aware Publish/Subscribe Framework for Parameterized Spatio-Textual Subscriptions

723  ChronoStream: Elastic Stateful Stream Computation in the Cloud

1504  Elaps: An Efficient Location-Aware Pub/Sub System

## Tang, Bo

207  Diversified Caching for Replicated Web Search Engines


## Tang, Nan

18  Proof Positive and Negative in Data Cleaning

## Tangwongsan, Kanat

149  Multicore Triangle Computations Without Tuning

## Tarek, Kareem

1585  SHAHED: A MapReduce-Based System for Querying and Visualizing Spatio-Temporal Satellite Data


*T continues on next page...*



## Telang, Aditya D.

999  Indexing and Matching Trajectories Under Inconsistent Sampling Rates


## Thian, Benji

1480  VISUAL: An HCI-inspired Simulator for Blending Visual Subgraph Query Construction and Processing


## Tian, Yuanyuan

1143  Dynamic Interaction Graphs with Probabilistic Edge Decay

## Tirthapura, Srikanta

243  Mining Maximal Cliques from an Uncertain Graph

## To, Hien


1404  PrivGeoCrowd: A Toolbox for Studying Private Spatial Crowdsourcing

## Tong, Yongxin


6  Cleaning Uncertain Data with a Noisy Crowd

735  Conflict-Aware Event-Participant Arrangement


## Torp, Kristian

1412  EcoSky: Reducing Vehicular Environmental Impact Through Eco-Routing

## Towsley, Don

939  A Tale of Three Graphs: Sampling Design on Hybrid Social-Affiliation Networks

## Tözün, Pınar

1530  How to Stop Under-Utilization and Love Multicores


## Tung, Anthony K.H.

113  PABIRS: A Data Access Middleware for Distributed File Systems



# U

U, Leong Hou

579  Quick-Motif: An Efficient and Scalable Framework for Exact Motif Discovery

Ukkonen, Antti

1155  Bump Hunting in the Dark: Local Discrepancy Maximization on Graphs


Ulusoy, Huseyin

78  Privacy-Aware Dynamic Feature Selection



# V

## Velegrakis, Yannis


1538  Goals in Social Media, Information Retrieval and Intelligent Agents

1573  Fine-Grained Controversy Detection in Wikipedia


## Verroios, Vasilis

219  Entity Resolution with Crowd Errors


## Vesdapunt, Norases

627  Identifying Users in Social Networks with Limited Information

## Vijayarajendran, Priya

1400  Advanced Analytics on SAP HANA: Churn Risk Scoring Using Call Network Analysis

## Voigt, Hannes

1460  Enjoy FRDM — Play with a Schema-Flexible RDBMS




# W



## Waguih, Dalia Attia

- 1440  AllegatorTrack: Combining and Reporting Results of Truth Discovery from Multi-Source Data

## Wang, Chaoyong

- 1448  PGWinFunc: Optimizing Window Aggregate Functions in PostgreSQL and its Application for Trajectory Data

## Wang, Haixun

- 495  Short Text Understanding Through Lexical-Semantic Analysis
- 1534  Inferencing in Information Extraction: Techniques and Applications

## Wang, Jianmin

- 30  Cleaning Structured Event Logs: A Graph Repair Approach

## Wang, Jianyong

- 1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream

## Wang, Jin

- 519  *Two Birds with One Stone*: An Efficient Hierarchical Framework for Top-k and Threshold-Based String Similarity Search

## Wang, Jiye

- 1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive

## Wang, Pinghui

- 939  A Tale of Three Graphs: Sampling Design on Hybrid Social-Affiliation Networks

*W continues on next page...*




## Wang, Shaoqing

1420  CROWN: A Context-Aware Recommender for Web News


## Wang, Ting

1316  FCCE: Highly Scalable Distributed Feature Collection and Correlation Engine for Low Latency Big Data Analytics


## Wang, Wei

1107  AP-Tree: Efficiently Support Continuous Spatial-Keyword Queries over Stream

## Wang, Xiang

1107  AP-Tree: Efficiently Support Continuous Spatial-Keyword Queries over Stream


## Wang, Xiaoling

1448  PGWinFunc: Optimizing Window Aggregate Functions in PostgreSQL and its Application for Trajectory Data


## Wang, Yan

351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs

## Wang, Yazhe

54  Preserving Privacy in Social Networks Against Connection Fingerprint Attacks

## Wang, Zhongyuan

495  Short Text Understanding Through Lexical-Semantic Analysis

## Watanabe, Yousuke

1292  AEDSMS: Automotive Embedded Data Stream Management System

## Wei, Mingzhu

783  INSURE: An Integrated Load Reduction Framework for XML Stream Processing

*W continues on next page...*



## Weidlich, Matthias

231  Result Selection and Summarization for Web Table Search

1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks

## Weidner, Martin

1259  Towards a Web-Scale Data Management Ecosystem Demonstrated by SAP HANA

## Weigold, Thomas

1328  Seamlessly Integrating Disk and Tape in a Multi-Tiered Distributed File System

## Wen, Zeyi

531  The Safest Path via Safe Zones

## Wiener, Janet L.

1516  Open Data Challenges at Facebook

## Wijaya, Tri Kurniawan

1488  SMART: A Tool for Analyzing and Reconciling Schema Matching Networks

## Winslett, Marianne

101  Chronos: An Elastic Parallel Framework for Stream Benchmark Generation and Simulation

113  PABIRS: A Data Access Middleware for Distributed File Systems

## Wong, Weng-Fai

1268  "Anti-Caching"-based Elastic Memory Management for Big Data

## Wu, Buwen

795  Scalable SPARQL Querying Using Path Partitioning

## Wu, Ke

651  False Rumors Detection on Sina Weibo by Propagation Structures


*W continues on next page...*



## Wu, Sai

113  PABIRS: A Data Access Middleware for Distributed File Systems


## Wu, Shanchan

1424  LearningAssistant: A Novel Learning Resource Recommendation System

## Wu, Shensen

1268  “Anti-Caching”-based Elastic Memory Management for Big Data

## Wu, Wei

1452  *CDR-To-MoVis*: Developing a *Mobility Visualization* System from CDR Data

## Wu, Yingjun

723  ChronoStream: Elastic Stateful Stream Computation in the Cloud

## Wu, Yubao

915  Finding Dense and Connected Subgraphs in Dual Networks





# X

Xia, Yubin

1268  “Anti-Caching”-based Elastic Memory Management for Big Data


Xiao, Xiaokui

831  Network Motif Discovery: A GPU Approach

Xiao, Zheng

1340  DualTable: A Hybrid Storage Model for Update Optimization in Hive

Xie, Wenlei


1143  Dynamic Interaction Graphs with Probabilistic Edge Decay

Xie, Xiaojing


1492  PIGEON: Progress Indicator for Subgraph Queries

Xie, Xing

831  Network Motif Discovery: A GPU Approach

975  Approximate Keyword Search in Semantic Trajectory Database

Xiong, Li

1035  Differentially Private Frequent Sequence Mining via Sampling-Based Candidate Pruning


Xu, Chuanfei

207  Diversified Caching for Replicated Web Search Engines

Xu, Guandong

1436  MARS: A Multi-Aspect Recommender System for Point-of-Interest


Xu, Jiajie

423  Interactive Top- $k$  Spatial Keyword Queries

*X continues on next page...*




## Xu, Jianliang


279  Answering Why-Not Questions on Spatial Keyword Top- $k$  Queries

339  Asymmetric Structure-Preserving Subgraph Queries for Large Graphs

## Xu, Pan

243  Mining Maximal Cliques from an Uncertain Graph

## Xu, Shengzhi

1035  Differentially Private Frequent Sequence Mining via Sampling-Based Candidate Pruning





# Y


## Yamaguchi, Akihiro

1292  AEDSMS: Automotive Embedded Data Stream Management System

## Yang, Bin

267  Finding Top- $k$  Local Users in Geo-Tagged Social Media Data

543  Personalized Route Recommendation Using Big Trajectory Data

1412  EcoSky: Reducing Vehicular Environmental Impact Through Eco-Routing

## Yang, Mohan

867  Optimizing Recursive Queries with Monotonic Aggregates in *DeALS*


## Yang, Song

651  False Rumors Detection on Sina Weibo by Propagation Structures

## Yap, Ghim-Eng

1452  *CDR-To-MoVis*: Developing a *Mobility Visualization* System from CDR Data

## Yi, Ke

173  Piecewise Linear Approximation of Streaming Time Series Data with Max-Error Guarantees

## Yi, Peipei

1492  PIGEON: Progress Indicator for Subgraph Queries

## Yiu, Man Lung

207  Diversified Caching for Replicated Web Search Engines

579  Quick-Motif: An Efficient and Scalable Framework for Exact Motif Discovery

*Y continues on next page...*



## Yu, Cong

1534  Inferencing in Information Extraction: Techniques and Applications


## Yu, Jeffrey Xu


927  On Random Walk Based Graph Sampling

## Yu, Philip S.

185  On Historical Diagnosis of Sensor Streams

## Yu, Ting


1023  Conservative or Liberal? Personalized Differential Privacy

1512  Interactive Preference-Aware Query Optimization


## Yu, Xiaohui


411  Meaningful Keyword Search in Relational Databases with Large and Complex Schema

## Yuan, Long

387  Diversified Top-K Clique Search

## Yuan, Nicholas Jing


963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach

975  Approximate Keyword Search in Semantic Trajectory Database


## Yuan, Pingpeng

795  Scalable SPARQL Querying Using Path Partitioning

## Yuan, Quan

675  SAR: A Sentiment-Aspect-Region Model for User Preference Analysis in Geo-Tagged Reviews


## Yuan, Xiaojie

375  Answering Regular Path Queries on Workflow Provenance




# Z

## Zait, Mohamed

1253  Oracle Database In-Memory: A Dual Format In-Memory Database

## Zampetakis, Stamatis


771  CliqueSquare: Flat Plans for Massively Parallel RDF Queries


1432  CliqueSquare in Action: Flat Plans for Massively Parallel RDF Queries

## Zaniolo, Carlo

867  Optimizing Recursive Queries with Monotonic Aggregates in *DeALS*

## Zeng, Kai

963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach

1059  Groupwise Analytics via Adaptive MapReduce

## Zhang, Chao

1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream

## Zhang, Chen Jason

6  Cleaning Uncertain Data with a Noisy Crowd


## Zhang, Dongxiang

1504  Elaps: An Efficient Location-Aware Pub/Sub System

## Zhang, Fan

1376  Growing the Charging Station Network for Electric Vehicles with Trajectory Data Analytics


## Zhang, Hao

1268  “Anti-Caching”-based Elastic Memory Management for Big Data

*Z continues on next page...*



## Zhang, Jinbo

- 1500  DAVINCI: Data-Driven Visual Interface Construction for Subgraph Search in Graph Databases

## Zhang, Rong

- 483  Dish Comment Summarization Based on Bilateral Topic Analysis

## Zhang, Rui

- 531  The Safest Path via Safe Zones

## Zhang, Wei

- 1561  STREAMCUBE: Hierarchical Spatio-Temporal Hashtag Clustering for Event Exploration over the Twitter Stream

## Zhang, Wenjie


- 387  Diversified Top-K Clique Search

- 1107  AP-Tree: Efficiently Support Continuous Spatial-Keyword Queries over Stream

## Zhang, Xiang

- 915  Finding Dense and Connected Subgraphs in Dual Networks

## Zhang, Ying

- 1107  AP-Tree: Efficiently Support Continuous Spatial-Keyword Queries over Stream





## Zhang, Yong

- 519  *Two Birds with One Stone: An Efficient Hierarchical Framework for Top-k and Threshold-Based String Similarity Search*

*Z continues on next page...*




## Zhang, Zhenjie

- 101  Chronos: An Elastic Parallel Framework for Stream Benchmark Generation and Simulation
- 113  PABIRS: A Data Access Middleware for Distributed File Systems
- 483  Dish Comment Summarization Based on Bilateral Topic Analysis
- 567  A General Graph-Based Model for Recommendation in Event-Based Social Networks

## Zhao, Dongyan

- 1508  A Graph-Based RDF Triple Store

## Zhao, Junzhou

- 939  A Tale of Three Graphs: Sampling Design on Hybrid Social-Affiliation Networks

## Zhao, Kaiqi

- 675  SAR: A Sentiment-Aspect-Region Model for User Preference Analysis in Geo-Tagged Reviews



## Zhao, Kankan

- 1420  CROWN: A Context-Aware Recommender for Web News

## Zhao, Kun

- 987  Bounded Quadrant System: Error-Bounded Trajectory Compression on the Go

## Zhao, Lei

- 66  Efficient Secure Similarity Computation on Encrypted Trajectory Data
- 351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs



## Zheng, Baihua

- 54  Preserving Privacy in Social Networks Against Connection Fingerprint Attacks







*Z continues on next page...*






## Zheng, Bolong

- 423  Interactive Top- $k$  Spatial Keyword Queries
- 975  Approximate Keyword Search in Semantic Trajectory Database


## Zheng, Kai

- 66  Efficient Secure Similarity Computation on Encrypted Trajectory Data
- 351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs
- 423  Interactive Top- $k$  Spatial Keyword Queries
- 495  Short Text Understanding Through Lexical-Semantic Analysis
- 963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach
- 975  Approximate Keyword Search in Semantic Trajectory Database

## Zhou, Aoying

- 101  Chronos: An Elastic Parallel Framework for Stream Benchmark Generation and Simulation
- 483  Dish Comment Summarization Based on Bilateral Topic Analysis
- 1448  PGWinFunc: Optimizing Window Aggregate Functions in PostgreSQL and its Application for Trajectory Data

## Zhou, Minqi

- 101  Chronos: An Elastic Parallel Framework for Stream Benchmark Generation and Simulation

## Zhou, Shuigeng

- 1492  PIGEON: Progress Indicator for Subgraph Queries

## Zhou, Xianke







- 113  PABIRS: A Data Access Middleware for Distributed File Systems

*Z continues on next page...*






## Zhou, Xiaofang

- 66  Efficient Secure Similarity Computation on Encrypted Trajectory Data
- 351  Multi-Constrained Graph Pattern Matching in Large-Scale Contextual Social Graphs
- 423  Interactive Top- $k$  Spatial Keyword Queries
- 495  Short Text Understanding Through Lexical-Semantic Analysis
- 963  Making Sense of Trajectory Data: A Partition-and-Summarization Approach
- 975  Approximate Keyword Search in Semantic Trajectory Database



## Zhou, Yongluan

- 795  Scalable SPARQL Querying Using Path Partitioning

## Zhu, Feida

- 1500  DAVINCI: Data-Driven Visual Interface Construction for Subgraph Search in Graph Databases

## Zhu, Kenny Q.

- 651  False Rumors Detection on Sina Weibo by Propagation Structures
- 675  SAR: A Sentiment-Aspect-Region Model for User Preference Analysis in Geo-Tagged Reviews

## Zhu, Xiaochen

- 30  Cleaning Structured Event Logs: A Graph Repair Approach

## Zhu, Xiaofeng

- 915  Finding Dense and Connected Subgraphs in Dual Networks

## Zhuang, Fuzhen

- 891  Online Frequent Episode Mining

*Z continues on next page...*




## Zou, Benyou

1420  CROWN: A Context-Aware Recommender for Web News


## Zou, Lei

1508  A Graph-Based RDF Triple Store

## Züfle, Andreas

1428  A Framework for Computation of Popular Paths from Crowdsourced Data

## Zukowski, Marcin

315  Smooth Scan: Statistics-Oblivious Access Paths

## Zwilling, Mike

1239  Evolving the Architecture of SQL Server for Modern Hardware Trends