




A

Abadi, Daniel J.


64  Query Optimization of Distributed Pattern Matching

Abbar, Sofiane


460  Ranking Item Features by Mining Online User-Item Interactions

Abedjan, Ziawasch

1036  Detecting Unique Column Combinations on Dynamic Data

1198  Profiling and Mining RDF Data with ProLOD++

Aberer, Karl

220  Pay-as-You-Go Reconciliation in Schema Matching Networks

Agarwal, Pankaj K.

748  Top- k Preferences in High Dimensions


Agrawal, Ankit


1072  SILVERBACK: Scalable Association Mining for Temporal Data in Columnar Probabilistic Databases


Agrawal, Divyakant

496  Pagrol: PArallel GRaph OLap Over Large-Scale Attributed Graphs

Ailamaki, Anastasia

1  Running with Scissors: Fast Queries on Just-in-Time Databases

688  ATraPos: Adaptive Transaction Processing on Hardware Islands

1000  OCTOPUS: Efficient Query Execution on Dynamic Mesh Datasets

A continues on next page...



Ajmera, Jitendra

436  Automatic Generation of Question Answer Pairs from Noisy Case Logs

Alarabi, Louai

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator


Alonso, Omar

1202  Kondenzer: Exploration and Visualization of Archived Social Media

Alvaro, Peter

52  Blazes: Coordination Analysis for Distributed Programs


Aly, Ahmed M.

1238  Mars: Real-Time Spatio-Temporal Queries on Microblogs


Amer-Yahia, Sihem

412  Exploiting Group Recommendation Functions for Flexible Preferences

Anastasiu, David C.

784  L2AP: Fast Cosine Similarity Search with Prefix L-2 Norm Bounds

Assent, Ira

88  Discriminative Features for Identifying and Interpreting Outliers




B

Bao, Jie


1096  Region Sampling and Estimation of GeoSocial Data with Dynamic Range Calibration

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator

Bao, Zhifeng

940  Breaking out of the MisMatch Trap


Basu Roy, Senjuti

412  Exploiting Group Recommendation Functions for Flexible Preferences

Bertino, Elisa

640  Practical k Nearest Neighbor Queries with Location Privacy

Bharathan, Vivek


1084  DBDesigner: A Customizable Physical Design Tool for Vertica Analytic Database

Billawala, Youssef

1168  Stock Trade Volume Prediction with Yahoo Finance User Browsing Behavior

Bodagala, Sreenath

1084  DBDesigner: A Customizable Physical Design Tool for Vertica Analytic Database

1108  The Vertica Query Optimizer: The Case for Specialized Query Optimizers

Boden, Brigitte

796  PHiDJ: Parallel Similarity Self-Join for High-Dimensional Vector Data with MapReduce

Böhlen, Michael

100  Memory-Efficient Centroid Decomposition for Long Time Series

B continues on next page...



Bordino, Ilaria

1168  Stock Trade Volume Prediction with Yahoo Finance User Browsing Behavior


Boysen, Mikkel

1194  Constructing Indoor Navigation Systems from Digital Building Information

Bruckner, Daniel

1258  Text and Structured Data Fusion in Data Tamer at Scale

Buyya, Rajkumar

808  MELODY-JOIN: Efficient Earth Mover's Distance Similarity Joins Using MapReduce



C

Cam, Hasan

820  Top-K Interesting Subgraph Discovery in Information Networks

Canahuate, Guadalupe


484  A Tunable Compression Framework for Bitmap Indices

Candan, K. Selçuk

388  Leveraging Metadata for Identifying Local, Robust Multi-Variate Temporal (RMT) Features

Cao, Caleb Chen

1174  C-DMr: Crowd-Powered Decision Maker for Real World Knapsack Problems

1182  CrowdCleaner: Data Cleaning for Multi-Version Data on the Web via Crowdsourcing


Cao, Lei

76  Scalable Distance-Based Outlier Detection Over High-Volume Data Streams

Cao, Yu

1274  HOPE: Iterative and Interactive Database Partitioning for OLTP Workloads


Cary, Ariel

1084  DBDesigner: A Customizable Physical Design Tool for Vertica Analytic Database


Castillo, Carlos

460  Ranking Item Features by Mining Online User-Item Interactions

Cerf, Loïc

448  Complete Discovery of High-Quality Patterns in Large Numerical Tensors

Cetindil, Inci

328  Efficient Instant-Fuzzy Search with Proximity Ranking


C continues on next page...



Cha, Youngchul

1156  Exploration of the Effect of Category Match Score in Search Advertising


Chaitanya, Arun

532  Decorrelation of User Defined Function Invocations in Queries


Chang, Hui Joe

1132  Efficient Support of XQuery Full Text in SQL/XML Enabled RDBMS

Cheema, Muhammad Aamir

760  SLICE: Reviving Regions-Based Pruning for Reverse k Nearest Neighbors Queries

Chen, Gang

40  R-Store: A Scalable Distributed System for Supporting Real-Time Analytics


Chen, Jian


808  MELODY-JOIN: Efficient Earth Mover's Distance Similarity Joins Using MapReduce

Chen, Lei


28  Towards Effective and Efficient Mining of Arbitrary Shaped Clusters

1144  CrowdPlanner: A Crowd-Based Route Recommendation System

1174  *C-DMr*: Crowd-Powered Decision Maker for Real World Knapsack Problems

1182  CrowdCleaner: Data Cleaning for Multi-Version Data on the Web via Crowdsourcing

Chen, Wei

1206  Trendspedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Chen, Zhao

1174  *C-DMr*: Crowd-Powered Decision Maker for Real World Knapsack Problems

C continues on next page...




Cheng, Hong

736  Random-Walk Domination in Large Graphs


Cheng, Reynold

724  Evaluating Multi-Way Joins Over Discounted Hitting Time


1186  iTag: Incentive-Based Tagging

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data

Cherniack, Mitch

1278  Devel-Op: An Optimizer Development Environment

Chiang, Fei

244  Continuous Data Cleaning

Chiew, Kevin

28  Towards Effective and Efficient Mining of Arbitrary Shaped Clusters

Chiu, David

484  A Tunable Compression Framework for Bitmap Indices

Cho, Junghoo

1156  Exploration of the Effect of Category Match Score in Search Advertising

Choi, Byron

628  Private Search on Key-Value Stores with Hierarchical Indexes

Choudhary, Alok

1072  SILVERBACK: Scalable Association Mining for Temporal Data in Columnar Probabilistic Databases


Christodoulakis, Christina

1230  VoidWiz: Resolving Incompleteness Using Network Effects

C continues on next page...



Chrysanthis, Panos K.

424  PAQO: Preference-Aware Query Optimization for Decentralized Database Systems

Chu, Xu

1222  RULEMINER: Data Quality Rules Discovery

Cohen, Sara

928  A General Algorithm for Subtree Similarity-Search


Conway, Neil

52  Blazes: Coordination Analysis for Distributed Programs




D

Dang, Xuan Hong

88  Discriminative Features for Identifying and Interpreting Outliers

Das, Gautam


412  Exploiting Group Recommendation Functions for Flexible Preferences

460  Ranking Item Features by Mining Online User-Item Interactions

Das Sarma, Anish

964  Crowd-Powered Find Algorithms

Dave, Jaimin

1084  DBDesigner: A Customizable Physical Design Tool for Vertica Analytic Database

1108  The Vertica Query Optimizer: The Case for Specialized Query Optimizers

de Haas, Christian

1194  Constructing Indoor Navigation Systems from Digital Building Information

Deng, Dong

340  MassJoin: A MapReduce-Based Method for Scalable String Similarity Joins

Deshpande, Amol

904  Subgraph Pattern Matching Over Uncertain Graphs with Identity Linkage Uncertainty



E

El Abbadi, Amr


496  Pagrol: PArallel GRaph OLap Over Large-Scale Attributed Graphs

Eldawy, Ahmed

1242  Pigeon: A Spatial MapReduce Language

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator

Elmasri, Ramez


1250  GQBE: Querying Knowledge Graphs by Example Entity Tuples

Elmehdwi, Yousef


664  Secure k -Nearest Neighbor Query Over Encrypted Data in Outsourced Environments

Elnikety, Sameh

172  Mercury: A Memory-Constrained Spatio-Temporal Real-Time Search on Microblogs

1238  Mars: Real-Time Spatio-Temporal Queries on Microblogs


Enrich, Tobias

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data

Ernst, Patrick

1254  KnowLife: A Knowledge Graph for Health and Life Sciences

Esmaelnezhad, Jamshid

328  Efficient Instant-Fuzzy Search with Proximity Ranking




F

Faloutsos, Christos

1230  VoidWiz: Resolving Incompleteness Using Network Effects

Fan, Ju

976  A Hybrid Machine-Crowdsourcing System for Matching Web Tables


Fan, Qi

496  Pagrol: PARallel GRaph OLap Over Large-Scale Attributed Graphs


Fan, Wenfei

184  Answering Graph Pattern Queries Using Views

Farnan, Nicholas L.

424  PAQO: Preference-Aware Query Optimization for Decentralized Database Systems

Fekete, Alan

700  Scalable Serializable Snapshot Isolation for Multicore Systems

Feng, Jianhua

340  MassJoin: A MapReduce-Based Method for Scalable String Similarity Joins

880  Effective Location Identification from Microblogs

Feng, Wei

856  We Can Learn Your #Hashtags: Connecting Tweets to Explicit Topics


Fries, Sergej

796  PHiDJ: Parallel Similarity Self-Join for High-Dimensional Vector Data with MapReduce



G

Gal, Avigdor

220  Pay-as-You-Go Reconciliation in Schema Matching Networks

Gamper, Johann

100  Memory-Efficient Centroid Decomposition for Long Time Series

Gao, Jing

820  Top-K Interesting Subgraph Discovery in Information Networks

Gao, Jun

544  GLog: A High Level Graph Analysis System Using MapReduce

556  Continuous Pattern Detection Over Billion-Edge Graph Using Distributed Framework

Gao, Lixin

316  Fast Top-K Path-Based Relevance Query on Massive Graphs


Gao, Yunjun

28  Towards Effective and Efficient Mining of Arbitrary Shaped Clusters


Garcia-Molina, Hector

964  Crowd-Powered Find Algorithms

Garofalakis, Minos

124  Continuous Fragmented Skylines Over Distributed Streams

Geerts, Floris

232  Mapping and Cleaning

Getoor, Lise

904  Subgraph Pattern Matching Over Uncertain Graphs with Identity Linkage Uncertainty

G continues on next page...



Ghinita, Gabriel

844  Large-Scale Frequent Subgraph Mining in MapReduce


Golab, Lukasz

1290  Data Stream Warehousing

Gómez Ferro, Daniel

676  Omid: Lock-Free Transactional Support for Distributed Data Stores

Gruetze, Toni

1198  Profiling and Mining RDF Data with ProLOD++

Gubanov, Michael

1258  Text and Structured Data Fusion in Data Tamer at Scale


Guo, Chenjuan

136  Stochastic Skyline Route Planning Under Time-Varying Uncertainty

Guo, Xiaoyan

1274  HOPE: Iterative and Interactive Database Partitioning for OLTP Workloads

Gupta, Mahesh

1250  GQBE: Querying Knowledge Graphs by Example Entity Tuples

Gupta, Manish

820  Top-K Interesting Subgraph Discovery in Information Networks

Gupta, Rajeev

1306  Distributed Execution of Continuous Queries


Gupta, Vani

652  Generating Private Synthetic Databases for Untrusted System Evaluation

G continues on next page...



Guravannavar, Ravindra

532  Decorrelation of User Defined Function Invocations in Queries

Güting, Ralf Hartmut

1190  Parallel SECONDO: A Practical System for Large-Scale Processing of Moving Objects

Guzun, Gheorghii

484  A Tunable Compression Framework for Bitmap Indices




H

Hacıgümüş, Hakan

712  Automatic Entity-Grouping for OLTP Workloads

1266  SAGE: A Logical and Physical Design Tool for Entity-Group Based New SQL Systems

Hakuta, Shizuya

952  XQuery Streaming by Forest Transducers


Halevy, Alon

964  Crowd-Powered Find Algorithms

Halverson, Alan

352  In-RDBMS Inverted Indexes Revisited

Han, Hyuck

700  Scalable Serializable Snapshot Isolation for Multicore Systems


Han, Jiawei

820  Top-K Interesting Subgraph Discovery in Information Networks

Hao, Shuang

340  MassJoin: A MapReduce-Based Method for Scalable String Similarity Joins

Harb, Boulos

1120  Near Neighbor Join

Hartig, Olaf

1286  Linked Data Query Processing

Hassan, Naeemul

112  Incremental Discovery of Prominent Situational Facts


H continues on next page...



He, Qinming


28  Towards Effective and Efficient Mining of Arbitrary Shaped Clusters

He, Wenbo


1012  An Efficient Sampling Method for Characterizing Points of Interests on Maps

He, Yuxiong

172  Mercury: A Memory-Constrained Spatio-Temporal Real-Time Search on Microblogs

1238  Mars: Real-Time Spatio-Temporal Queries on Microblogs

Heinis, Thomas


1000  OCTOPUS: Efficient Query Execution on Dynamic Mesh Datasets

Hellerstein, Joseph M.


52  Blazes: Coordination Analysis for Distributed Programs

Hsiung, Wang-Pin


712  Automatic Entity-Grouping for OLTP Workloads

1266  SAGE: A Logical and Physical Design Tool for Entity-Group Based New SQL Systems


Hu, Haibo

628  Private Search on Key-Value Stores with Hierarchical Indexes

Hu, Jun

880  Effective Location Identification from Microblogs

Hu, Xin

1214  Outsourcing Multi-Version Key-Value Stores with Verifiable Data Freshness

Hu, Zhirui

280  Head, Modifier, and Constraint Detection in Short Texts

H continues on next page...



Huang, Hao


28  Towards Effective and Efficient Mining of Arbitrary Shaped Clusters

Huang, Jiamin


1144  CrowdPlanner: A Crowd-Based Route Recommendation System

1178  A Crowd-Based Route Recommendation System — CrowdPlanner

Huang, Jiewen

64  Query Optimization of Distributed Pattern Matching


Huang, Jin

808  MELODY-JOIN: Efficient Earth Mover's Distance Similarity Joins Using MapReduce

Huang, Xin

736  Random-Walk Domination in Large Graphs

Hwang, Jeong-Hyon

196  Efficient Top- k Closeness Centrality Search



I

Idreos, Stratos

1262  dbTouch in Action — Database Kernels for Touch-Based Data Exploration


Ilyas, Ihab F.

1222  RULEMINER: Data Quality Rules Discovery

Islam, Md. Saiful

268  Keyword-Based Correlated Network Computation Over Large Social Media

Iwasaki, Hideya

952  XQuery Streaming by Forest Transducers




J

Jacobsen, Hans-Arno


364  Adaptive Parallel Compressed Event Matching

1162  CaSSanDra: An SSD Boosted Key-Value Store


Jagadish, H.V.

940  Breaking out of the MisMatch Trap


Jain, Rohit

1282  Guaranteed Authenticity and Integrity of Data from Untrusted Servers

Jayachandran, Prasanth

472  Distributed and Interactive Cube Exploration

Jayaram, Nandish


1250  GQBE: Querying Knowledge Graphs by Example Entity Tuples

Jensen, Christian S.

136  Stochastic Skyline Route Planning Under Time-Varying Uncertainty

148  Scalable Top-*k* Spatio-Temporal Term Querying

Jentzsch, Anja

1198  Profiling and Mining RDF Data with ProLOD++




Jermaine, Christopher M.

520  History-Aware Query Optimization with Materialized Intermediate Views

Jeung, Hoyoung

1144  CrowdPlanner: A Crowd-Based Route Recommendation System

J continues on next page...

**Jiang, Di**400  Personalized Query Suggestion with Diversity Awareness**Jiang, Wei**664  Secure k -Nearest Neighbor Query Over Encrypted Data in Outsourced Environments**Jin, Tan**892  Efficient and Accurate Query Evaluation on Uncertain Graphs via Recursive Stratified Sampling**Johnson, Theodore**1290  Data Stream Warehousing**Joshi, Sachindra**436  Automatic Generation of Question Answer Pairs from Noisy Case Logs**Jovanovic, Petar**1060  Engine Independence for Logical Analytic Flows**Jung, Hyungsoo**700  Scalable Serializable Snapshot Isolation for Multicore Systems**Jung, Kyomin**292  LinkSCAN*: Overlapping Community Detection Using the Link-Space Transformation**Junqueira, Flavio**676  Omid: Lock-Free Transactional Support for Distributed Data Stores




K

Kamat, Niranjana

472  Distributed and Interactive Cube Exploration

Kang, Wei


868  Interactive Hierarchical Tag Clouds for Summarizing Spatiotemporal Social Contents

1206  Trendspedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Kao, Ben

724  Evaluating Multi-Way Joins Over Discounted Hitting Time

Karypis, George

784  L2AP: Fast Cosine Similarity Search with Prefix L-2 Norm Bounds


Katib, Anas

1270  A Tool for Internet-Scale Cardinality Estimation of XPath Queries Over Distributed Semistructured Data

Kaul, Manohar

136  Stochastic Skyline Route Planning Under Time-Varying Uncertainty


Kefaloukos, Pimin Konstantin


1024  Declarative Cartography: In-Database Map Generalization of Geospatial Datasets


Kelly, Ivan

676  Omid: Lock-Free Transactional Support for Distributed Data Stores

Kemper, Alfons

580  Exploiting Hardware Transactional Memory in Main-Memory Databases

592  Locality-Sensitive Operators for Parallel Main-Memory Database Clusters

1310  Main-Memory Database Systems

K continues on next page...


[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](#)

**Kersten, Martin**

508  Waste Not... Efficient Co-Processing of Relational Data

Khan, Arijit

1250  GQBE: Querying Knowledge Graphs by Example Entity Tuples

Khandelwal, Kartikay

1202  Kondenzer: Exploration and Visualization of Archived Social Media


Khayati, Mourad

100  Memory-Efficient Centroid Decomposition for Long Time Series

Khemmarat, Samamon

316  Fast Top-K Path-Based Relevance Query on Massive Graphs

Kim, Taewoo

328  Efficient Instant-Fuzzy Search with Proximity Ranking

Kimmig, Angelika

904  Subgraph Pattern Matching Over Uncertain Graphs with Identity Linkage Uncertainty


Kllapi, Herald

1120  Near Neighbor Join


Kondreddi, Sarath Kumar

988  Combining Information Extraction and Human Computing for Crowdsourced Knowledge Acquisition

Korn, Flip

1226  iCoDA: Interactive and Exploratory Data Completeness Analysis


Kourtellis, Nicolas

1168  Stock Trade Volume Prediction with Yahoo Finance User Browsing Behavior

K continues on next page...



Kriegel, Hans-Peter

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data


Kwon, Sejeong

292  LinkSCAN*: Overlapping Community Detection Using the Link-Space Transformation



L

Labouseur, Alan G.

196  Efficient Top- k Closeness Centrality Search

Lakshmanan, Laks V.S.

3  Incremental Cluster Evolution Tracking from Highly Dynamic Network Data


Lamb, Andrew

1108  The Vertica Query Optimizer: The Case for Specialized Query Optimizers

Laptev, Nikolay

1168  Stock Trade Volume Prediction with Yahoo Finance User Browsing Behavior

Lee, Adam J.

424  PAQO: Preference-Aware Query Optimization for Decentralized Database Systems

Lee, Jae-Gil

292  LinkSCAN*: Overlapping Community Detection Using the Link-Space Transformation

Lee, Kang-Won

832  Cloud Service Placement via Subgraph Matching


Lee, Pei

3  Incremental Cluster Evolution Tracking from Highly Dynamic Network Data


Lei, Siyu

1186  iTag: Incentive-Based Tagging

Leis, Viktor

580  Exploiting Hardware Transactional Memory in Main-Memory Databases

Leung, Kenneth Wai-Ting

400  Personalized Query Suggestion with Diversity Awareness


L continues on next page...

[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](#)




Li, Chen


328  Efficient Instant-Fuzzy Search with Proximity Ranking

Li, Chengkai

112  Incremental Discovery of Prominent Situational Facts

1250  GQBE: Querying Knowledge Graphs by Example Entity Tuples


Li, Feng

40  R-Store: A Scalable Distributed System for Supporting Real-Time Analytics

Li, Guoliang

340  MassJoin: A MapReduce-Based Method for Scalable String Similarity Joins

880  Effective Location Identification from Microblogs


940  Breaking out of the MisMatch Trap

Li, Jianxin


268  Keyword-Based Correlated Network Computation Over Large Social Media


Li, Rong-Hua

736  Random-Walk Domination in Large Graphs

892  Efficient and Accurate Query Evaluation on Uncertain Graphs via Recursive Stratified Sampling

Li, Xinyu

868  Interactive Hierarchical Tag Clouds for Summarizing Spatiotemporal Social Contents

1206  Trendspectra: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Li, Yanhua

1096  Region Sampling and Estimation of GeoSocial Data with Dynamic Range Calibration


L continues on next page...



Li, Yatao

1182  CrowdCleaner: Data Cleaning for Multi-Version Data on the Web via Crowdsourcing

Liarou, Erietta

688  ATraPos: Adaptive Transaction Processing on Hardware Islands

1262  dbTouch in Action — Database Kernels for Touch-Based Data Exploration

Lim, Sungsu


292  LinkSCAN*: Overlapping Community Detection Using the Link-Space Transformation

Lin, Wenqing


844  Large-Scale Frequent Subgraph Mining in MapReduce

Lin, Xuemin


304  Fast Incremental SimRank on Link-Evolving Graphs

760  SLICE: Reviving Regions-Based Pruning for Reverse k Nearest Neighbors Queries


Lindstrom, Peter


616  Optimal Hierarchical Layouts for Cache-Oblivious Search Trees

Ling, Tok Wang

940  Breaking out of the MisMatch Trap

Liu, Bin

712  Automatic Entity-Grouping for OLTP Workloads

1266  SAGE: A Logical and Physical Design Tool for Entity-Group Based New SQL Systems


Liu, Chengfei

268  Keyword-Based Correlated Network Computation Over Large Social Media


L continues on next page...



Liu, Ling

1214  Outsourcing Multi-Version Key-Value Stores with Verifiable Data Freshness


Liu, Ruilin

1226  iCoDA: Interactive and Exploratory Data Completeness Analysis

Liu, Tianyu

1178  A Crowd-Based Route Recommendation System — CrowdPlanner

Liu, Xue

1012  An Efficient Sampling Method for Characterizing Points of Interests on Maps

Liu, Zhen Hua

1132  Efficient Support of XQuery Full Text in SQL/XML Enabled RDBMS

Lu, Hua

1194  Constructing Indoor Navigation Systems from Digital Building Information

Lu, Jiamin

1190  Parallel SECONDO: A Practical System for Large-Scale Processing of Moving Objects

Lu, Meiyu

976  A Hybrid Machine-Crowdsourcing System for Matching Web Tables

Lu, Wentian

652  Generating Private Synthetic Databases for Untrusted System Evaluation

Lu, Ying

1132  Efficient Support of XQuery Full Text in SQL/XML Enabled RDBMS




M

Madden, Samuel

604  Rethinking Main Memory OLTP Recovery

Magdy, Amr

172  Mercury: A Memory-Constrained Spatio-Temporal Real-Time Search on Microblogs

1238  Mars: Real-Time Spatio-Temporal Queries on Microblogs

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator


Maier, David

52  Blazes: Coordination Analysis for Distributed Programs

Malviya, Nirmesh

604  Rethinking Main Memory OLTP Recovery


Mamoulis, Nikos

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data

Manegold, Stefan

508  Waste Not... Efficient Co-Processing of Relational Data


Maneth, Sebastian

952  XQuery Streaming by Forest Transducers

Maniu, Silviu

1186  iTag: Incentive-Based Tagging


Mao, Rui

892  Efficient and Accurate Query Evaluation on Uncertain Graphs via Recursive Stratified Sampling


M continues on next page...




Markram, Henry


1000  OCTOPUS: Efficient Query Execution on Dynamic Mesh Datasets

Mecca, Giansalvatore

232  Mapping and Cleaning

1218  IQ-METER — An Evaluation Tool for Data-Transformation Systems

Meira Jr., Wagner

448  Complete Discovery of High-Quality Patterns in Large Numerical Tensors

Meng, Cynthia

1254  KnowLife: A Knowledge Graph for Health and Life Sciences

Menon, Prashanth

1162  CaSSanDra: An SSD Boosted Key-Value Store

Miele, Antonio

1234  ADaPT: Automatic Data Personalization Based on Contextual Preferences

Miklau, Gerome

652  Generating Private Synthetic Databases for Untrusted System Evaluation


Miklós, Zoltán

220  Pay-as-You-Go Reconciliation in Schema Matching Networks

Milios, Evangelos E.

3  Incremental Cluster Evolution Tracking from Highly Dynamic Network Data

Miller, Renée J.

244  Continuous Data Cleaning

1230  VoidWiz: Resolving Incompleteness Using Network Effects

M continues on next page...



Mittal, Amol


436  Automatic Generation of Question Answer Pairs from Noisy Case Logs

Mo, Luyi

1186  iTag: Incentive-Based Tagging

Mokbel, Mohamed F.

172  Mercury: A Memory-Constrained Spatio-Temporal Real-Time Search on Microblogs

1238  Mars: Real-Time Spatio-Temporal Queries on Microblogs


1242  Pigeon: A Spatial MapReduce Language

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator

Moustafa, Walaa Eldin

904  Subgraph Pattern Matching Over Uncertain Graphs with Identity Linkage Uncertainty


Mühlbauer, Tobias

592  Locality-Sensitive Operators for Parallel Main-Memory Database Clusters



N

Nakano, Keisuke


952  XQuery Streaming by Forest Transducers

Nandi, Arnab

472  Distributed and Interactive Cube Exploration

Nath, Suman


172  Mercury: A Memory-Constrained Spatio-Temporal Real-Time Search on Microblogs


1238  Mars: Real-Time Spatio-Temporal Queries on Microblogs

Naughton, Jeffrey F.


352  In-RDBMS Inverted Indexes Revisited


Naumann, Felix


1036  Detecting Unique Column Combinations on Dynamic Data

1198  Profiling and Mining RDF Data with ProLOD++

Neumann, Thomas

580  Exploiting Hardware Transactional Memory in Main-Memory Databases


592  Locality-Sensitive Operators for Parallel Main-Memory Database Clusters

1310  Main-Memory Database Systems

Ng, Raymond T.

88  Discriminative Features for Identifying and Interpreting Outliers

Ng, Wilfred

400  Personalized Query Suggestion with Diversity Awareness

N continues on next page...



Nguyen, Quoc Viet Hung

220  Pay-as-You-Go Reconciliation in Schema Matching Networks


Nguyen, Thanh Tam

220  Pay-as-You-Go Reconciliation in Schema Matching Networks





O

Olsen Jr., Paul W.


196  Efficient Top- k Closeness Centrality Search

Ooi, Beng Chin

40  R-Store: A Scalable Distributed System for Supporting Real-Time Analytics

976  A Hybrid Machine-Crowdsourcing System for Matching Web Tables

Or, Nerya

928  A General Algorithm for Subtree Similarity-Search

Özsu, M. Tamer


40  R-Store: A Scalable Distributed System for Supporting Real-Time Analytics

1286  Linked Data Query Processing




P

Palsetia, Diana

1072  SILVERBACK: Scalable Association Mining for Temporal Data in Columnar Probabilistic Databases


Papaemmanouil, Olga

1278  Devel-Op: An Optimizer Development Environment

Papapetrou, Odysseas

124  Continuous Fragmented Skylines Over Distributed Streams

Papotti, Paolo

232  Mapping and Cleaning


1218  IQ-METER — An Evaluation Tool for Data-Transformation Systems

1222  RULEMINER: Data Quality Rules Discovery

Parameswaran, Aditya

964  Crowd-Powered Find Algorithms


Park, SeongJae

700  Scalable Serializable Snapshot Isolation for Multicore Systems


Paulet, Russell

640  Practical k Nearest Neighbor Queries with Location Privacy

Pei, Kexin

628  Private Search on Key-Value Stores with Hierarchical Indexes


Peng, Peng

772  Geometry Approach for k-Regret Query


P continues on next page...




Peng, Zhibo

1278  Devel-Op: An Optimizer Development Environment

Perez, Luis L.

520  History-Aware Query Optimization with Materialized Intermediate Views

Pietzuch, Peter

1214  Outsourcing Multi-Version Key-Value Stores with Verifiable Data Freshness

Pilvinyte, Aiste

1194  Constructing Indoor Navigation Systems from Digital Building Information


Pirk, Holger

508  Waste Not... Efficient Co-Processing of Relational Data


Po, Oliver

712  Automatic Entity-Grouping for OLTP Workloads

Porobic, Danica

688  ATraPos: Adaptive Transaction Processing on Hardware Islands

Prabhakar, Sunil

1282  Guaranteed Authenticity and Integrity of Data from Untrusted Servers




Q

Qarabaqi, Bahar

916  User-Driven Refinement of Imprecise Queries

Qin, Lu

208  Contract & Expand: I/O Efficient SCCs Computing

Quiané-Ruiz, Jorge-Arnulfo

1036  Detecting Unique Column Combinations on Dynamic Data

Quintarelli, Elisa

1234  ADaPT: Automatic Data Personalization Based on Contextual Preferences



R

Rabl, Tilmann

1162  CaSSanDra: An SSD Boosted Key-Value Store

Rabosio, Emanuele

1234  ADaPT: Automatic Data Personalization Based on Contextual Preferences

Rae, Ian

352  In-RDBMS Inverted Indexes Revisited

Raghavendra, Ramya

832  Cloud Service Placement via Subgraph Matching


Rahman, Habibur

460  Ranking Item Features by Mining Online User-Item Interactions

Rajan, Deepak

616  Optimal Hierarchical Layouts for Cache-Oblivious Search Trees


Ramachandra, Karthik

532  Decorrelation of User Defined Function Invocations in Queries

Ramamritham, Krithi

1306  Distributed Execution of Continuous Queries

Rao, Praveen

1270  A Tool for Internet-Scale Cardinality Estimation of XPath Queries Over Distributed Semistructured Data

Ravishankar, China V.

256  On Masking Topical Intent in Keyword Search


R continues on next page...




Reed, Benjamin

676  Omid: Lock-Free Transactional Support for Distributed Data Stores

Reiser, Angelika

592  Locality-Sensitive Operators for Parallel Main-Memory Database Clusters

Renz, Matthias

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data


Riedewald, Mirek

916  User-Driven Refinement of Imprecise Queries

Rödiger, Wolf

592  Locality-Sensitive Operators for Parallel Main-Memory Database Clusters

Röhm, Uwe

700  Scalable Serializable Snapshot Isolation for Multicore Systems

Rundensteiner, Elke A.

76  Scalable Distance-Based Outlier Detection Over High-Volume Data Streams

Ryu, Seungwoo

292  LinkSCAN*: Overlapping Community Detection Using the Link-Space Transformation



S

Sadoghi, Mohammad


364  Adaptive Parallel Compressed Event Matching

1162  CaSSanDra: An SSD Boosted Key-Value Store

Saha, Barna

1294  Data Quality: The Other Face of Big Data


Sailer, Reiner


1214  Outsourcing Multi-Version Key-Value Stores with Verifiable Data Freshness

Samanthula, Bharath K.

664  Secure k -Nearest Neighbor Query Over Encrypted Data in Outsourced Environments

Santoro, Donatello

232  Mapping and Cleaning

1218  IQ-METER — An Evaluation Tool for Data-Transformation Systems


Sapino, Maria Luisa

388  Leveraging Metadata for Identifying Local, Robust Multi-Variate Temporal (RMT) Features


Sarwat, Mohamed

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator

Sawin, Jason

484  A Tunable Compression Framework for Bitmap Indices

Schubert, Erich

88  Discriminative Features for Identifying and Interpreting Outliers

S continues on next page...



Schürmann, Felix

1000  OCTOPUS: Efficient Query Execution on Dynamic Mesh Datasets


Seidl, Thomas

796  PhiDJ: Parallel Similarity Self-Join for High-Dimensional Vector Data with MapReduce

Shang, Shuo

136  Stochastic Skyline Route Planning Under Time-Varying Uncertainty


Shao, Bin

568  How to Partition a Billion-Node Graph


Sharaf, Mohamed A.

1210  AQUAS: A Quality-Aware Scheduler for NoSQL Data Stores

Sheth, Amit

2  Transforming Big Data into Smart Data: Deriving Value via Harnessing Volume, Variety, and Velocity Using Semantic Techniques and Technologies


Shrinivas, Lakshmikant

1108  The Vertica Query Optimizer: The Case for Specialized Query Optimizers


Sidlauskas, Darius

148  Scalable Top- k Spatio-Temporal Term Querying

Simhadri, Varun

532  Decorrelation of User Defined Function Invocations in Queries

Simitsis, Alkis

1060  Engine Independence for Logical Analytic Flows

S continues on next page...

**Singh, Ambuj K.**

160  Nearest Keyword Set Search in Multi-Dimensional Datasets

832  Cloud Service Placement via Subgraph Matching


Singh, Vishwakarma

160  Nearest Keyword Set Search in Multi-Dimensional Datasets

Siu, Amy

1254  KnowLife: A Knowledge Graph for Health and Life Sciences


Skovsgaard, Anders

148  Scalable Top- k Spatio-Temporal Term Querying

Slavov, Vasil

1270  A Tool for Internet-Scale Cardinality Estimation of XPath Queries Over Distributed Semistructured Data

Song, Qiyue

1206  Trendspedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Song, Shaoxu

376  Matching Heterogeneous Events with Patterns

Srivastava, Divesh

1294  Data Quality: The Other Face of Big Data

Srivatsa, Mudhakar

832  Cloud Service Placement via Subgraph Matching

Steiner, Moritz

1096  Region Sampling and Estimation of GeoSocial Data with Dynamic Range Calibration


S continues on next page...



Stepien, Grzegorz

796  PHiDJ: Parallel Similarity Self-Join for High-Dimensional Vector Data with MapReduce

Stonebraker, Michael

604  Rethinking Main Memory OLTP Recovery

1258  Text and Structured Data Fusion in Data Tamer at Scale

Su, Han

1144  CrowdPlanner: A Crowd-Based Route Recommendation System

1178  A Crowd-Based Route Recommendation System — CrowdPlanner

Su, Jianwen

1048  Modeling Data for Business Processes

Sudarshan, S.

532  Decorrelation of User Defined Function Invocations in Queries

Sultana, Afroza

112  Incremental Discovery of Prominent Situational Facts


Sun, Jiaguang

376  Matching Heterogeneous Events with Patterns

Sun, Yutian

1048  Modeling Data for Business Processes

Szlichta, Jaroslaw

244  Continuous Data Cleaning




T

Tan, Kian-Lee

496  Pagrol: PArallel GRaph OLap Over Large-Scale Attributed Graphs

880  Effective Location Identification from Microblogs


Tan, Wang-Chiew

976  A Hybrid Machine-Crowdsourcing System for Matching Web Tables

Tanca, Letizia


1234  ADaPT: Automatic Data Personalization Based on Contextual Preferences

Tang, Yuzhe


1214  Outsourcing Multi-Version Key-Value Stores with Verifiable Data Freshness

Tatemura, Junichi

712  Automatic Entity-Grouping for OLTP Workloads

1266  SAGE: A Logical and Physical Design Tool for Entity-Group Based New SQL Systems

Tauheed, Farhan

1000  OCTOPUS: Efficient Query Execution on Dynamic Mesh Datasets

Thirumuruganathan, Saravanan

412  Exploiting Group Recommendation Functions for Flexible Preferences

460  Ranking Item Features by Mining Online User-Item Interactions

Todd, Stephen

1274  HOPE: Iterative and Interactive Database Partitioning for OLTP Workloads


Tong, Yongxin

1182  CrowdCleaner: Data Cleaning for Multi-Version Data on the Web via Crowdsourcing

T continues on next page...




Tözün, Pınar

688  ATraPos: Adaptive Transaction Processing on Hardware Islands

Trajcevski, Goce

1072  SILVERBACK: Scalable Association Mining for Temporal Data in Columnar Probabilistic Databases

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data


Tran, Nga


1108  The Vertica Query Optimizer: The Case for Specialized Query Optimizers

Triantafillou, Peter


988  Combining Information Extraction and Human Computing for Crowdsourced Knowledge Acquisition

Tung, Anthony K.H.

868  Interactive Hierarchical Tag Clouds for Summarizing Spatiotemporal Social Contents

1206  Trendspedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web


Tunga, Karthik

472  Distributed and Interactive Cube Exploration



U


Unterbrunner, Philipp

592  Locality-Sensitive Operators for Parallel Main-Memory Database Clusters



V


Varadarajan, Ramakrishna

1084  DBDesigner: A Customizable Physical Design Tool for Vertica Analytic Database


Varadharajan, Vijay

640  Practical k Nearest Neighbor Queries with Location Privacy


Vaz Salles, Marcos

1024  Declarative Cartography: In-Database Map Generalization of Geospatial Datasets


Venkatraman, Kartik

64  Query Optimization of Distributed Pattern Matching


Verma, Ashish

436  Automatic Generation of Question Answer Pairs from Noisy Case Logs


Viglas, Stratis D.

1298  Just-in-Time Compilation for SQL Query Processing

Volkovs, Maksims

244  Continuous Data Cleaning


Vosecky, Jan

400  Personalized Query Suggestion with Diversity Awareness



W

Wang, Guan

1226  iCoDA: Interactive and Exploratory Data Completeness Analysis

Wang, Haixun

280  Head, Modifier, and Constraint Detection in Short Texts

568  How to Partition a Billion-Node Graph

Wang, Haozhou

1178  A Crowd-Based Route Recommendation System — CrowdPlanner


Wang, Huiju

496  Pagrol: PArallel GRaph OLap Over Large-Scale Attributed Graphs


Wang, Jianmin

376  Matching Heterogeneous Events with Patterns

Wang, Jiannan

340  MassJoin: A MapReduce-Based Method for Scalable String Similarity Joins

Wang, Jianyong

856  We Can Learn Your #Hashtags: Connecting Tweets to Explicit Topics


Wang, Jiayuan

76  Scalable Distance-Based Outlier Detection Over High-Volume Data Streams

Wang, Limin

1096  Region Sampling and Estimation of GeoSocial Data with Dynamic Range Calibration


Wang, Lu

568  How to Partition a Billion-Node Graph


W continues on next page...



Wang, Peng

256  On Masking Topical Intent in Keyword Search


Wang, Pinghui

1012  An Efficient Sampling Method for Characterizing Points of Interests on Maps


Wang, Qingyang

76  Scalable Distance-Based Outlier Detection Over High-Volume Data Streams

Wang, Ting

1214  Outsourcing Multi-Version Key-Value Stores with Verifiable Data Freshness

Wang, Wendy Hui

1226  iCoDA: Interactive and Exploratory Data Completeness Analysis

Wang, Xiaolan

388  Leveraging Metadata for Identifying Local, Robust Multi-Variate Temporal (RMT) Features

Wang, Xin

184  Answering Graph Pattern Queries Using Views

Wang, Zhengkui

496  Pagrol: PArallel GRaph OLap Over Large-Scale Attributed Graphs

Wang, Zhongyuan

280  Head, Modifier, and Constraint Detection in Short Texts

Waytas, Ethan

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator



Weidlich, Matthias

220  Pay-as-You-Go Reconciliation in Schema Matching Networks

W continues on next page...




Weikum, Gerhard

- 988  Combining Information Extraction and Human Computing for Crowdsourced Knowledge Acquisition
- 1254  KnowLife: A Knowledge Graph for Health and Life Sciences


Weisberg, Ariel

- 604  Rethinking Main Memory OLTP Recovery

Wilkinson, Kevin

- 1060  Engine Independence for Logical Analytic Flows

Wong, Raymond Chi-Wing

- 772  Geometry Approach for k-Regret Query

Wu, Budan

- 1048  Modeling Data for Business Processes

Wu, Yinghui

- 184  Answering Graph Pattern Queries Using Views



X

Xia, Fan

1210  AQUAS: A Quality-Aware Scheduler for NoSQL Data Stores

Xia, Leihao

1174  *C-DMr*: Crowd-Powered Decision Maker for Real World Knapsack Problems

Xiao, Xiaokui

844  Large-Scale Frequent Subgraph Mining in MapReduce

Xiao, Yanghua

568  How to Partition a Billion-Node Graph

Xie, Xike

1194  Constructing Indoor Navigation Systems from Digital Building Information

Xie, Yusheng

1072  SILVERBACK: Scalable Association Mining for Temporal Data in Columnar Probabilistic Databases


Xu, Chen

1210  AQUAS: A Quality-Aware Scheduler for NoSQL Data Stores

Xu, Jianliang

628  Private Search on Key-Value Stores with Hierarchical Indexes


Xu, Xizhong

628  Private Search on Key-Value Stores with Hierarchical Indexes



Y


Yabandeh, Maysam

676  Omid: Lock-Free Transactional Support for Distributed Data Stores

Yackel, Steven

1246  A Demonstration of MNTG — A Web-Based Road Network Traffic Generator


Yan, Guanhua

15  Finding Common Ground Among Experts' Opinions on Data Clustering: With Applications in Malware Analysis


Yan, Tak

1156  Exploration of the Effect of Category Match Score in Search Advertising

Yan, Xifeng

820  Top-K Interesting Subgraph Discovery in Information Networks

832  Cloud Service Placement via Subgraph Matching

1250  GQBE: Querying Knowledge Graphs by Example Entity Tuples

Yang, Bin

136  Stochastic Skyline Route Planning Under Time-Varying Uncertainty

Yang, Di

76  Scalable Distance-Based Outlier Detection Over High-Volume Data Streams

Yang, Jian

1048  Modeling Data for Business Processes

Y continues on next page...




Yang, Jun


112  Incremental Discovery of Prominent Situational Facts

748  Top- k Preferences in High Dimensions

Yang, Shiyu

760  SLICE: Reviving Regions-Based Pruning for Reverse k Nearest Neighbors Queries


Yang, Xuan S.

1186  iTag: Incentive-Based Tagging

Ye, Yin

1222  RULEMINER: Data Quality Rules Discovery

Yeom, Heon Y.

700  Scalable Serializable Snapshot Isolation for Multicore Systems

Yi, Xun


640  Practical k Nearest Neighbor Queries with Location Privacy


Yu, Albert

748  Top- k Preferences in High Dimensions

Yu, Cong

112  Incremental Discovery of Prominent Situational Facts






412  Exploiting Group Recommendation Functions for Flexible Preferences

1120  Near Neighbor Join

Y continues on next page...




Yu, Jeffrey Xu

- 208  Contract & Expand: I/O Efficient SCCs Computing
- 544  GLog: A High Level Graph Analysis System Using MapReduce
- 556  Continuous Pattern Detection Over Billion-Edge Graph Using Distributed Framework
- 736  Random-Walk Domination in Large Graphs
- 892  Efficient and Accurate Query Evaluation on Uncertain Graphs via Recursive Stratified Sampling

Yu, Philip S.

- 376  Matching Heterogeneous Events with Patterns

Yu, Ting

- 424  PAQO: Preference-Aware Query Optimization for Decentralized Database Systems

Yu, Weiren

- 304  Fast Incremental SimRank on Link-Evolving Graphs

Yu, Yanwei

- 76  Scalable Distance-Based Outlier Detection Over High-Volume Data Streams


Yuan, Jian

- 1156  Exploration of the Effect of Category Match Score in Search Advertising




Z


Zachariassen, Martin

1024  Declarative Cartography: In-Database Map Generalization of Geospatial Datasets


Zeng, Yong

940  Breaking out of the MisMatch Trap

Zhang, Chao

1206  Trendsmedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Zhang, Chen Jason

1182  CrowdCleaner: Data Cleaning for Multi-Version Data on the Web via Crowdsourcing

Zhang, Meihui

976  A Hybrid Machine-Crowdsourcing System for Matching Web Tables


Zhang, Rui

808  MELODY-JOIN: Efficient Earth Mover's Distance Similarity Joins Using MapReduce


Zhang, Wangda

724  Evaluating Multi-Way Joins Over Discounted Hitting Time


Zhang, Wenjie

304  Fast Incremental SimRank on Link-Evolving Graphs

Zhang, Ying

760  SLICE: Reviving Regions-Based Pruning for Reverse k Nearest Neighbors Queries



Zhang, Zhiwei

208  Contract & Expand: I/O Efficient SCCs Computing



Z continues on next page...



Zhao, Feng

- 868  Interactive Hierarchical Tag Clouds for Summarizing Spatiotemporal Social Contents
- 1206  Trendspedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Zheng, Kai

- 1144  CrowdPlanner: A Crowd-Based Route Recommendation System
- 1178  A Crowd-Based Route Recommendation System — CrowdPlanner



Zhou, Aoying

- 1210  AQUAS: A Quality-Aware Scheduler for NoSQL Data Stores



Zhou, Baoyao

- 1274  HOPE: Iterative and Interactive Database Partitioning for OLTP Workloads

Zhou, Chang

- 544  GLog: A High Level Graph Analysis System Using MapReduce
- 556  Continuous Pattern Detection Over Billion-Edge Graph Using Distributed Framework


Zhou, Jiashuai

- 544  GLog: A High Level Graph Analysis System Using MapReduce
- 556  Continuous Pattern Detection Over Billion-Edge Graph Using Distributed Framework


Zhou, Minqi

- 1210  AQUAS: A Quality-Aware Scheduler for NoSQL Data Stores

Zhou, Shuigeng

- 628  Private Search on Key-Value Stores with Hierarchical Indexes

Zhou, Xiajuan

- 1206  Trendspedia: An Internet Observatory for Analyzing and Visualizing the Evolving Web

Z continues on next page...



Zhou, Xiaofang

1144  CrowdPlanner: A Crowd-Based Route Recommendation System

1178  A Crowd-Based Route Recommendation System — CrowdPlanner


Zhu, Ting

1096  Region Sampling and Estimation of GeoSocial Data with Dynamic Range Calibration

Zhu, Xiaochen

376  Matching Heterogeneous Events with Patterns


Zimek, Arthur

88  Discriminative Features for Identifying and Interpreting Outliers

Zong, Bo

832  Cloud Service Placement via Subgraph Matching

Züfle, Andreas

1302  Managing Uncertainty in Spatial and Spatio-Temporal Data