

2014 IEEE International Congress on Big Data

(BigData Congress 2014)

**Anchorage, Alaska, USA
27 June – 2 July 2014**



**IEEE Catalog Number: CFP14SEV-POD
ISBN: 978-1-4799-5058-4**

2014 IEEE International Congress on Big Data

BigData Congress 2014

Table of Contents

Message from the Organizing Committee.....	xvii
Technical Program Committee.....	xviii
External Reviewers.....	xxiii
Satellite Sessions Organizing Committees.....	xxiv
IEEE Computer Society Technical Committee on Services Computing.....	xxviii
Service Society.....	xxix

Session 1

BigData Research Session 1 - BigData Analytics Techniques

GraphLens: Mining Enterprise Storage Workloads Using Graph Analytics	1
<i>Yang Zhou, Sangeetha Seshadri, Lawrence Chiu, and Ling Liu</i>	
FSM-H: Frequent Subgraph Mining Algorithm in Hadoop	9
<i>Mansurul A. Bhuiyan and Mohammad Al Hasan</i>	
Rectangle Counting in Large Bipartite Graphs	17
<i>Jia Wang, Ada Wai-Chee Fu, and James Cheng</i>	

BigData Research Session 2 - MapReduce Model

A Parallel Spatial Co-location Mining Algorithm Based on MapReduce	25
<i>Jin Soung Yoo, Douglas Boulware, and David Kimmey</i>	
Energy-Aware Scheduling of MapReduce Jobs	32
<i>Lena Mashayekhy, Mahyar Movahed Nejad, Daniel Grosu, Dajun Lu, and Weisong Shi</i>	
Vigiles: Fine-Grained Access Control for MapReduce Systems	40
<i>Huseyin Ulusoy, Murat Kantarcioglu, Erman Pattuk, and Kevin Hamlen</i>	

BigData Research Session 3 - BigData Security

Denial-of-Service Threat to Hadoop/YARN Clusters with Multi-tenancy	48
<i>Jingwei Huang, David M. Nicol, and Roy H. Campbell</i>	
A Big Data Architecture for Large Scale Security Monitoring	56
<i>Samuel Marchal, Xiuyan Jiang, Radu State, and Thomas Engel</i>	
Contextual Anomaly Detection in Big Sensor Data	64
<i>Michael A. Hayes and Miriam A.M. Capretz</i>	

BigData Research Session 4 - BigData Performance Issues

High-Performance Spatial Query Processing on Big Taxi Trip Data Using GPGPUs	72
<i>Jianting Zhang, Simin You, and Le Gruenwald</i>	
A Compatible LZMA ORC-Based Optimization for High Performance Big Data Load	80
<i>Liping Zhang, Qi Chen, and Kai Miao</i>	

BigData Research Session 5 - BigData Analytics

Storing a Collection of Differentially Compressed Files Recursively	88
<i>Angelos Molfetas, Anthony Wirth, and Justin Zobel</i>	
XDB - A Novel Database Architecture for Data Analytics as a Service	96
<i>Carsten Binnig, Abdallah Salama, Erfan Zamanian, Harald Kornmayer, Sven Listing, and Alexander C. Mueller</i>	
DeltaDB: A Scalable Database Design for Time-Varying Schema-Free Data	104
<i>Peter Ivie and Douglas Thain</i>	

BigData Research Session 6 - MapReduce Framework

Practising Scalable Graph Similarity Joins in MapReduce	112
<i>Yifan Chen, Xiang Zhao, Bin Ge, Chuan Xiao, and Chi-Hung Chi</i>	
Configuring a MapReduce Framework for Performance-Heterogeneous Clusters	120
<i>Jessica Hartog, Renan Delvalle, Madhusudhan Govindaraju, and Michael J. Lewis</i>	
Marimba: A Framework for Making MapReduce Jobs Incremental	128
<i>Johannes Schildgen, Thomas Jörg, Manuel Hoffmann, and Stefan Deßloch</i>	

BigData Research Session 7 - BigData Services

Money on the Move: Big Data of Bank Card Transactions as the New Proxy for Human Mobility Patterns and Regional Delineation. The Case of Residents and Foreign Visitors in Spain	136
<i>Stanislav Sobolevsky, Izabela Sitko, Remi Tachet Des Combes, Bartosz Hawelka, Juan Murillo Arias, and Carlo Ratti</i>	
Characterizing and Modeling Package Dynamics in Express Shipping Service Network	144
<i>Xu Tan, Yuanchao Shu, Xie Lu, Peng Cheng, and Jiming Chen</i>	
Dmodel: Online Taxicab Demand Model from Big Sensor Data in a Roving Sensor Network	152
<i>Desheng Zhang, Tian He, Shan Lin, Sirajum Munir, and John A. Stankovic</i>	

BigData Research Session 8 - Distributed BigData Services

Distributed Maximal Clique Computation	160
<i>Yanyan Xu, James Cheng, Ada Wai-Chee Fu, and Yingyi Bu</i>	
A Processing Pipeline for Cassandra Datasets Based on Hadoop Streaming	168
<i>E. Dede, B. Sendir, P. Kuzlu, J. Weachock, M. Govindaraju, and L. Ramakrishnan</i>	
AuDy: Automatic Dynamic Least-Weight Balancing for Stream Workloads Scalability	176
<i>Pedro Martins, Maryam Abbasi, and Pedro Furtado</i>	

BigData Research Session 9 - BigData Analytics Network

SiLK: A Tool Suite for Unsampled Network Flow Analysis at Scale	184
<i>Mark Thomas, Leigh Metcalf, Jonathan Spring, Paul Krystosek, and Katherine Prevost</i>	
Using Inter-file Similarity to Improve Intra-file Compression	192
<i>Angelos Molfetas, Anthony Wirth, and Justin Zobel</i>	

BigData Research Session 10 - BigData Management Model

A Cloud-Computing Local Histogram Construction Algorithm for Big Image Data	200
<i>Chung-Chih Cheng, Fan-Chieh Cheng, Po-Hsiung Lin, and Shih-Chia Huang</i>	
Scalable State Management for Scientific Applications in the Cloud	204
<i>Tonglin Li, Ioan Raicu, and Lavanya Ramakrishnan</i>	
T-PICE: Twitter Personality Based Influential Communities Extraction System	212
<i>Eleanna Kafeza, Andreas Kanavos, Christos Makris, and Pantelis Vikatos</i>	

A Holistic Framework for Big Scientific Data Management	220
<i>Verena Kantere</i>	

BigData Research Session 11 - Mexico City Satellite Session

Data Base Analysis Using a Compact Data Set	227
<i>Angel Fernando Kuri-Morales</i>	
The Babel File System	234
<i>Moisés Quezada Naquid, Ricardo Marcelín Jiménez, and Jose Luis González Compeán</i>	

BigData Research Session 12 - Coimbra Satellite Session

Using Multimedia Ontologies for Automatic Image Annotation and Classification	242
<i>Antonio M. Rinaldi</i>	
Real-Time Integration of Building Energy Data	250
<i>Diogo Anjos, Paulo Carreira, and Alexandre P. Francisco</i>	

BigData Research Session 13 - Taipei Satellite Session

Provenance as a Service: A Data-centric Approach for Real-Time Monitoring	258
<i>Rafat Hammad and Ching-Seh Wu</i>	
Big Data Collection Gateway for Vision-Based Smart Meter Reading Network	266
<i>Ching-Han Chen, Ching-Yi Chen, Chih-Hsien Hsia, and Guan-Xin Wu</i>	
Triangulation-Based Singer Identification for Duet Music Data Indexing	270
<i>Wei-Ho Tsai and Cin-Hao Ma</i>	
Speech and Singing Discrimination for Audio Data Indexing	276
<i>Wei-Ho Tsai and Cin-Hao Ma</i>	
Temporal Event Tracing on Big Healthcare Data Analytics	281
<i>Chin-Ho Lin, Liang-Cheng Huang, Seng-Cho T. Chou, Chih-Ho Liu, Han-Fang Cheng, and I-Jen Chiang</i>	
The Application of Cartesian-Join of Bloom Filters to Supporting Membership Query of Multidimensional Data	288
<i>Zhu Wang, Tiejian Luo, Guandong Xu, and Xiang Wang</i>	
Federated MapReduce to Transparently Run Applications on Multicluster Environment	296
<i>Chun-Yu Wang, Tzu-Li Tai, Shu Jui-Shing, Chang Jyh-Biau, and Shieh Ce-Kuen</i>	
How to Send a Self-Destructing Email: A Method of Self-Destructing Email System	304
<i>Xiao Fu, Zhijian Wang, Hao Wu, Jia-qi Yang, and Zi-zhao Wang</i>	

Association Rule Mining of Personal Hobbies in Social Networks	310
<i>Xiaoqing Yu, Huanhuan Liu, Jianhua Shi, Jenq-Neng Hwang, Wanggen Wan, and Jing Lu</i>	

BigData Research Session 14 - Data Processing

Reducing the Search Space for Big Data Mining for Interesting Patterns from Uncertain Data	315
<i>Carson Kai-Sang Leung, Richard Kyle MacKinnon, and Fan Jiang</i>	
Lethe: Cluster-Based Indexing for Secure Multi-user Search	323
<i>Eirini C. Micheli, Giorgos Margaritis, and Stergios V. Anastasiadis</i>	
Adaptive Indexing for Distributed Array Processing	331
<i>Yifeng Geng, Xiaomeng Huang, and Guangwen Yang</i>	

BigData Research Session 15 - Shenzhen Satellite Session

Hydrological Time Series Anomaly Mining Based on Symbolization and Distance Measure	339
<i>Dingsheng Wan, Yan Xiao, Pengcheng Zhang, Jun Feng, Yuelong Zhu, and Qian Liu</i>	
Enterprise Relationship Network: Build Foundation for Social Business	347
<i>Liqiang Wang, Shijun Liu, Li Pan, Lei Wu, and Xiangxu Meng</i>	
Towards Big Data Bayesian Network Learning - An Ensemble Learning Based Approach	355
<i>Yan Tang, Yu Wang, Kendra M.L. Cooper, and Ling Li</i>	
A Study on Sentiment Computing and Classification of Sina Weibo with Word2vec	358
<i>Bai Xue, Chen Fu, and Zhan Shaobin</i>	
A Study on Trend Prediction in Sina Weibo Community	364
<i>Chen Fu, Zhan Shaobin, and Shi Guangjun</i>	
MapCheckReduce: An Improved MapReduce Computing Model for Imprecise Applications	366
<i>Changjian Wang, Yuxing Peng, Mingxing Tang, Dongsheng Li, Shanshan Li, and Pengfei You</i>	
Three-Level Views of the Web Service Network: An Empirical Study Based on ProgrammableWeb	374
<i>Saixia Lyu, Jianxun Liu, Mingdong Tang, Guosheng Kang, Buqing Cao, and Yucong Duan</i>	
Predicting Wikipedia Editor's Editing Interest Based on Factor Graph Model	382
<i>Haisu Zhang, Sheng Zhang, Zhaolin Wu, Liwei Huang, and Yutao Ma</i>	

On Retrieving Moving Objects Gathering Patterns from Trajectory Data via Spatio-temporal Graph	390
<i>Junming Zhang, Jinglin Li, Shangguang Wang, Zhihan Liu, Quan Yuan, and Fangchun Yang</i>	

BigData Industry and Application Session 1 - Distributed Data Model

BigData Oriented Open Scalable Relational Data Model	398
<i>Zhiyun Zheng, Zhimeng Du, Lun Li, and Yike Guo</i>	
ADAM - A Database and Information Retrieval System for Big Multimedia Collections	406
<i>Ivan Giangreco, Ihab Al Kabary, and Heiko Schuldt</i>	
Building a Massive Stream Computing Platform for Flexible Applications	414
<i>Tianjian Chen, Zhengrui Man, Hao Li, Xin Sun, Raymond K. Wong, and Zhiwei Yu</i>	

BigData Industry and Application Session 2 - Cloud-Based BigData Management

OpenStack Café: A Novel Time-Based User-centric Resource Management Framework in the Cloud	422
<i>Alan Yu Shyang Tan, Ryan Kok Leong Ko, and Grace P.Y. Ng</i>	
Trend Analysis of Time Series Data Using Data Mining Techniques	430
<i>Arpit Baheti and Durga Toshniwal</i>	
Parallel POD Compression of Time-Varying Big Datasets Using m-Swap on the K Computer	438
<i>Chongke Bi, Kenji Ono, and Lu Yang</i>	

BigData Industry and Application Session 3 - BigData Management System

A Key-Value Based Application Platform for Enterprise Big Data	446
<i>Bo Hu, Yutao Ma, Liang-Jie Zhang, Jiake Shi, and Jiayan Zhong</i>	
Does Social Media Big Data Make the World Smaller? An Exploratory Analysis of Keyword-Hashtag Networks	454
<i>Ahmed Abdeen Hamed and Xindong Wu</i>	
Data Visualization in Educational Datasets Using a Rule-Based Inference System	462
<i>Aniruddha Desai, Muaz Mian, David Hazel, Ankur Teredesai, and Gregory Benner</i>	

BigData Industry and Application Session 4 - BigData Analytics Method

- A Rule-Based Temporal Analysis Method for Online Health Analytics and Its Application for Real-Time Detection of Neonatal Spells470
Anirudh Thommandram, J. Mikael Eklund, Carolyn McGregor, James Edward Pugh, and Andrew G. James
- A Discriminant Framework for Detecting Similar Scientific Research Projects Based on Big Data Mining478
Shanqing Li, Lirong Song, and Hui Zhao

BigData Industry and Application Session 5 - BigData Analytics Algorithm

- Distributed Implementation of Latent Rating Pattern Sharing Based Cross-domain Recommender System Approach482
Anil Kumar, Vikas Kapur, Apangshu Saha, Rajeev Kumar Gupta, Arun Singh, Santanu Chaudhury, and Sumeet Agarwal
- Auto-configuration System and Algorithms for Big Data-Enabled Internet-of-Things Platforms490
Apostolos Papageorgiou, Manuel Zahn, and Ernö Kovacs
- DualIso: An Algorithm for Subgraph Pattern Matching on Very Large Labeled Graphs498
Matthew Saltz, Ayushi Jain, Abhishek Kothari, Arash Fard, John A. Miller, and Lakshmish Ramaswamy

BigData Industry and Application Session 6 - BigData Mining

- An Autonomic Mediation Framework for Complex Physical Environments506
Philippe Lalanda and Catherine Hamon
- FSBD: A Framework for Scheduling of Big Data Mining in Cloud Computing514
Leila Ismail, Mohammad M. Masud, and Latifur Khan
- Towards a Semantic Extract-Transform-Load (ETL) Framework for Big Data Integration522
Srividya K. Bansal

BigData Industry and Application Session 7 - NoSQL

- Big Data Analytics for Predictive Manufacturing Control - A Case Study from Process Industry530
Julian Krumeich, Sven Jacobi, Dirk Werth, and Peter Loos
- BigObject Store: In-Place Computing for Interactive Analytics538
Yi-Cheng Huang, Wenwey Hseush, Yu-Chun Lai, and Michael Fong

Scalable XPath Evaluation on Large-Scale Continuously Evolving XML Repositories	546
<i>Phani Rohit Mullangi, Gowtham Penematsa, and Lakshmith Ramaswamy</i>	

BigData Industry and Application Session 8 - Big Spatial Data

Bridging the Gap between Spatial Data Sources and Mashup Applications	554
<i>Wei Zhou, Chi-Hung Chi, Can Wang, Raymond Wong, and Chen Ding</i>	
A Combinatorial Approach to Multidimensional Scaling	562
<i>Jorge Alencar, Tibérius O. Bonates, and Carlile Lavor</i>	

BigData Industry and Application Session 9 - Hadoop

Big R: Large-Scale Analytics on Hadoop Using R	570
<i>Oscar D. Lara Yejas, Weiqiang Zhuang, and Adarsh Pannu</i>	
Diagnosing Virtualized Hadoop Performance from Benchmark Results: An Exploratory Study	578
<i>Jun Fan, Xinhui Li, Chi Harold Liu, Jeffrey Buell, Gavin Lu, and Luke Lu</i>	
Performance Implications of SSDs in Virtualized Hadoop Clusters	586
<i>Sungyong Ahn, Sangkyu Park, Jae-Ki Hong, and Wooseok Chang</i>	

BigData Industry and Application Session 10 - BigData Privacy

We Know Where You Are Tweeting From: Assigning a Type of Place to Tweets Using Natural Language Processing and Random Forests	594
<i>Abdulkareem Alsudais, Gondy Leroy, and Anthony Corso</i>	
Making Big Data, Privacy, and Anonymization Work Together in the Enterprise: Experiences and Issues	601
<i>Jeff Sedayao, Rahul Bhardwaj, and Nakul Gorade</i>	

BigData Industry and Application Session 11 - Cloud-Based BigData Storage

CloudJet4BigData: Streamlining Big Data via an Accelerated Socket Interface	608
<i>Frank Zhigang Wang, Theo Dimitrakos, Na Helian, Sining Wu, Ling Li, and Rodric Yates</i>	
Resource Provisioning with QoS in Cloud Storage	616
<i>Wei-Chih Huang, Chuan-Ming Liu, and Chuan-Chi Lai</i>	
Data Sharing for Cloud Computing Platforms	621
<i>Marwan Sabbouh, Kenneth McCracken, and Geoff Cooney</i>	

BigData Industry and Application Session 12 - BigData Analytics Architecture

Towards a Set Theoretical Approach to Big Data Analytics	629
<i>Raghava Rao Mukkamala, Abid Hussain, and Ravi Vatrapu</i>	
Efficient String Sorting on Multi - and Many-Core Architectures	637
<i>Aleksandr Drozd, Miquel Pericàs, and Satoshi Matsuoka</i>	
A Tale of Two Data-Intensive Paradigms: Applications, Abstractions, and Architectures	645
<i>Shantenu Jha, Judy Qiu, Andre Luckow, Pradeep Mantha, and Geoffrey C. Fox</i>	

BigData Industry and Application Session 13 - NoSQL

ODBAPI: A Unified REST API for Relational and NoSQL Data Stores	653
<i>Rami Sellami, Sami Bhiri, and Bruno Defude</i>	
Terms Mining in Document-Based NoSQL: Response to Unstructured Data	661
<i>Richard K. Lomotey and Ralph Deters</i>	
Spatio-temporal Pseudo Relevance Feedback for Large-Scale and Heterogeneous Scientific Repositories	669
<i>Shin'ichi Takeuchi, Yuhei Akahoshi, Bun Theang Ong, Komei Sugiura, and Koji Zettsu</i>	
An Intelligent System for Forecasting the Trend of Consumed Electricity	677
<i>Hang Yang, Huajun Chen, Cai Yuan, and Fang Lianhang</i>	

BigData Industry and Application Session 14 - NoSQL

Tactical Clinical Text Mining for Improved Patient Characterization	683
<i>Heather Champion, Nick Pizzi, and Raja Krishnamoorthy</i>	
An Effective Integrated Method for Learning Big Imbalanced Data	691
<i>Mojgan Ghanavati, Raymond K. Wong, Fang Chen, Yang Wang, and Chang-Shing Perng</i>	

BigData Industry and Application Session 15 - Big Social Data

Multifaceted Visualisation of Annotated Social Media Data	699
<i>Sanat Kumar Bista, Surya Nepal, and Cécile Paris</i>	

BigData Industry and Application Session 16 - Graph Analytics

Enumerating Maximal Bicliques from a Large Graph Using MapReduce	707
<i>Arko Provo Mukherjee and Srikanta Tirthapura</i>	
LightGraph: Lighten Communication in Distributed Graph-Parallel Processing	717
<i>Yue Zhao, Kenji Yoshigoe, Mengjun Xie, Suijian Zhou, Remzi Seker, and Jiang Bian</i>	
RABID: A Distributed Parallel R for Large Datasets	725
<i>Hao Lin, Shuo Yang, and Samuel P. Midkiff</i>	

BigData Industry and Application Session 17 - Cloud-Based BigData Service

Trace-Based Analysis and Prediction of Cloud Computing User Behavior Using the Fractal Modeling Technique	733
<i>Shuang Chen, Mahboobeh Ghorbani, Yanzhi Wang, Paul Bogdan, and Massoud Pedram</i>	
Inherent Replica Inconsistency in Cassandra	740
<i>Xiangdong Huang, Jianmin Wang, Jian Bai, Guiguang Ding, and Mingsheng Long</i>	
Towards Emulation of Large Scale Complex Network Workloads on Graph Databases with XGDBench	748
<i>Miyuru Dayarathna and Toyotaro Suzumura</i>	

BigData Industry and Application Session 18 - Data Cleaning Model and Performance

BigData Work-in-Progress Session 1 - BigData Processing

Policy Design Based on Risk at Big Data Era: Case Study of Privacy Invasion in South Korea	756
<i>Hyejung Moon, Hyun Suk Cho, Seo Hwa Jeong, and Jangho Park</i>	
Knowledge Discovery from Big Data for Intrusion Detection Using LDA	760
<i>Jingwei Huang, Zbigniew Kalbarczyk, and David M. Nicol</i>	
Big Data Security and Privacy Issues in Healthcare	762
<i>Harsh Kupwade Patil and Ravi Seshadri</i>	
GrandLand Traffic Data Processing Platform	766
<i>Xingcan Cui, Zhen Dong, Liwei Lin, Renyong Song, and Xiaohui Yu</i>	
A Memory-Based Continuous Query Index for Stream Processing	768
<i>Cuiwen Xiong, Peng Zhang, Yan Li, Shipeng Zhang, Qingyun Liu, and Jianlong Tan</i>	

BigData Work-in-Progress Session 2 - BigData Analytics

Ranking Keyword Search Results with Query Logs	770
<i>Jing Zhou, Xiaohui Yu, Yang Liu, and Ziqiang Yu</i>	
Constructing an Issue Network from the Perspective of Common R&D Keywords	772
<i>Namgyu Kim, William Wong Xiu Shun, Jieun Kim, Kee-Young Kwahk, Seungryul Jeong, and Hyunchul Ahn</i>	
Agriculture Data for All - Integrated Tools for Agriculture Data Integration, Analytics, and Sharing	774
<i>Jarek Nabrzyski, Cheng Liu, Charles Vardeman, Sandra Gesing, and Milan Budhatoki</i>	
Asynchronous Service Analysis of Cloud DVR DataCenter	776
<i>Zhen Zhao</i>	
Work in Progress - In-Memory Analysis for Healthcare Big Data	778
<i>Muaz Mian, Ankur Teredesai, David Hazel, Sreenivasulu Pokuri, and Krishna Uppala</i>	
Large-Scale Heterogeneous Program Retrieval through Frequent Pattern Discovery and Feature Correlation Analysis	780
<i>Bo Liu, Liang Wu, Qiuxiang Dong, and Yuanchun Zhou</i>	

BigData Work-in-Progress Session 3 - BigData Management Framework

Towards Efficient KNN Joins on Data Streams	782
<i>Chong Yang, Xiaohui Yu, and Yang Liu</i>	
CouchFS: A High-Performance File System for Large Data Sets	784
<i>Fangzhou Yao and Roy H. Campbell</i>	
A Computational Framework for Integrating and Retrieving Biodiversity Data on a Large Scale	786
<i>Daniel L. Da Silva, Pedro L.P. Corrêa, Silvio L. Stanzani, Paulo André Filipak, and Andreiwid Sheffer Corrêa</i>	
Versioning Complex Data	788
<i>Matt Macduff, Benno Lee, and Sherman Beus</i>	
On the Way to Big Data Applications in Industrial Computed Tomography	792
<i>Alexander Ditter, Dietmar Fey, Tobias Schön, and Steven Oeckl</i>	

BigData Work-in-Progress Session 4 - BigData Decision Support System

Towards a Novel Model for Distributed Big Data Service Composition Using Functional Graph Matching	794
<i>Carlos R. Rivero and Hasan M. Jamil</i>	
Enrichment Patterns for Big Data	796
<i>Kerrie Holley, Gandhi Sivakumar, and Kalapriya Kannan</i>	
YCSB and TPC-H: Big Data and Decision Support Benchmarks	800
<i>Melyssa Barata, Jorge Bernardino, and Pedro Furtado</i>	
Understanding Student Behaviors in Online Classroom: Data Scientific Approach	802
<i>Jongbok Byun, Diane Pennington, Jorge Cardenas, Srabasti Dutta, and Jeral Kirwan</i>	
Scientific Data Infrastructure for Sustainability Science Mobile Applications	804
<i>Katherine G. Herbert, Emily Hill, Jerry Alan Falls, Joseph O. Ajala, Richard T. Boniface, and Paul W. Cushman</i>	
Really Opened Government Data: A Collaborative Transparency at Sight	806
<i>Andreiwid Sheffer Corrêa, Pedro L.P. Corrêa, Daniel L. Silva, and Flávio Soares Corrêa Da Silva</i>	
Classifying and Aggregating Context Attributes for Business Service Requests - No 'One-Size-Fits-All'	808
<i>Rebecca Copeland and Noel Crespi</i>	
Big Data Infrastructure for analyzing data generated by Wireless Sensor Networks	816
<i>Lidice Garcia Rios and Jos´e Alberto Incera Diguez</i>	
Automated Trading with Machine Learning on Big Data	824
<i>Dymitr Ruta</i>	
Author Index	831