1. Record Nr. UNINA9910484246603321

Titolo	On the Move to Meaningful Internet Systems. OTM 2017 Conferences : Confederated International Conferences: CoopIS, C&TC, and ODBASE 2017, Rhodes, Greece, October 23-27, 2017, Proceedings, Part II / / edited by Hervé Panetto, Christophe Debruyne, Walid Gaaloul, Mike Papazoglou, Adrian Paschke, Claudio Agostino Ardagna, Robert Meersman
Pubbl/distr/stampa	Cham : , : Springer International Publishing : , : Imprint : Springer, , 2017
ISBN	3-319-69459-6
Edizione	[1st ed. 2017.]
Descrizione fisica	1 online resource (XXVII, 515 p. 152 illus.)
Collana	Programming and Software Engineering ; ; 10574
Disciplina	004.678
Soggetti	Artificial intelligence Software engineering Application software Database management Data mining Mathematical logic Artificial Intelligence Software Engineering Information Systems Applications (incl. Internet) Database Management Data Mining and Knowledge Discovery Mathematical Logic and Formal Languages
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Intro -- General Co-chairs and Editors' Message for OnTheMove 2017 -- Organization -- OnTheMove 2017 Keynotes -- Pragmatic Semantics at Web Scale -- Evolution of the Industrial Internet of Things: Preparing for Change -- Contents - Part II -- Contents - Part I -- Cloud and Trusted Computing (C&TC) 2017 -- En -- Property Preserving Encryption in NoSQL Wide Column Stores -- 1 Introduction -- 2 Related

Work -- 3 Background -- 3.1 The Data Model of Wide Column Stores -- 3.2 Property-Preserving Encryption -- 4 Selecting Practical PPE Schemes -- 4.1 Deterministic Encryption -- 4.2 Order-Preserving Encryption -- 4.3 Searchable Encryption -- 5 Data Management on Server Side -- 5.1 An Onion Layer Model for WCSs -- 5.2 Querying the Encrypted Data -- 6 Evaluation -- 6.1 Implementation -- 6.2 Performance -- 7 Security -- 8 Extensions and Future Work -- 9 Conclusion -- References -- Towards a JSON-Based Fast Policy Evaluation Framework -- 1 Introduction -- 2 Fundamentals of JACPoL -- 3 Performance Evaluation -- 4 Conclusion -- References -- Gibbon: An Availability Evaluation Framework for Distributed Databases -- 1 Introduction -- 2 Background -- 2.1 Distributed Database Systems -- 2.2 Availability for DDBMSs -- 2.3 Cloud Infrastructure -- 3 Failure Impact Analysis -- 3.1 Replication for Availability -- 3.2 Failures and Recovery -- 4 Availability Metrics -- 4.1 Input Parameters -- 4.2 Output Metrics -- 5 Availability Evaluation -- 5.1 Evaluation Process -- 5.2 DB Gibbon Algorithm -- 6 Architecture -- 7 Discussion -- 7.1 Evaluation Scenario: MongoDB -- 7.2 Evaluation Results -- 8 Related Work -- 9 Conclusion and Future Work -- References -- Locality-Aware GC Optimisations for Big Data Workloads -- 1 Introduction -- 2 Background -- 2.1 NoSQL Databases and Object Locality -- 2.2 Garbage Collection Algorithms and Heap Layouts.
3 Gang Promotion in Heap Management -- 3.1 Identifying Relevant Object-Graphs -- 3.2 Gang GC on a Large Heap -- 4 Deployment of Gang Promotion on Hotspot JVM -- 4.1 Java Runtime Instrumentation -- 4.2 LAG1 --- Locality-Aware Extension of G1 -- 5 Evaluation -- 5.1 Evaluation Setup -- 5.2 Program Locality in System -- 5.3 Application Behaviour -- 6 Related Work -- 7 Conclusion -- References -- FairCloud: Truthful Cloud Scheduling with Continuous and Combinatorial Auctions -- 1 Introduction -- 2 Related Work -- 3 Architecture -- 3.1 Architecture Design Requirements -- 3.2 Bid Representation Language -- 3.3 Entities and Events -- 3.4 Auctioneer Algorithms -- 4 Implementation -- 5 Evaluation -- 6 Conclusion -- References -- A Novel WebGIS-Based Situational Awareness Platform for Trustworthy Big Data Integration and Analytics in Mobility Context -- 1 Introduction -- 2 The MAGNIFIER WebGIS-Based Situational Awareness Platform -- 3 The MAGNIFIER Technological Architecture -- 4 PRACTIONIST DSS -- 5 Conclusion and Future Works -- References -- On the Verification of Software Vulnerabilities During Static Code Analysis Using Data Mining Techniques -- Abstract -- 1 Introduction -- 2 Related Work -- 3 Approach -- 4 Conclusion, Discussion, and Future Work -- References -- International Conference on Ontologies, DataBases, and Applications of Semantics (ODBASE) 2017 -- En -- Linked Data and Ontology Reference Model for Infectious Disease Reporting Systems -- 1 Introduction -- 2 Goal of Semantics Usage -- 3 Semantic Reference Model -- 3.1 Semantifying Relational Data -- 3.2 Consistent Model -- 4 Geo Political Ontology -- 5 Relevant Work -- 6 Evaluation Through Competency Questions -- 6.1 Use Case -- 6.2 Competency Questions -- References -- PFSgeo: Preference-Enriched Faceted Search for Geographical Data -- 1 Introduction -- 2 Related Work and Novelty.
3 The PFSgeo Approach -- 3.1 Extensions of the Language of PFS -- 3.2 Extensions of the System Hippalus -- 3.3 Algorithmic Perspective -- 4 Performance Evaluation -- 4.1 Testing Scalability -- 4.2 Analysis of the Performance -- 4.3 Comparative Experimental Results -- 5 Evaluation with Users -- 5.1 Preparation -- 6 Concluding Remarks -- References -- Speeding up Publication of Linked Data Using Data Chunking in LinkedPipes ETL -- 1 Introduction -- 2 Data Chunking --

3 Data Chunking in LinkedPipes ETL -- 3.1 SPARQL Endpoint Extractor -- 3.2 Tabular Extractor and Files to RDF -- 3.3 SPARQL Transformers -- 3.4 Merger Component and Loaders -- 4 Datasets and Pipelines for Experiments -- 4.1 LISTID - List of Czech Business Entity IDs -- 4.2 ARES - Czech Business Registry -- 4.3 MONITOR - Budget Information of All Levels of Czech State Administration -- 4.4 RUIAN - Registry of Territorial Identification, Addresses and Real Estate -- 5 Experiments -- 5.1 LISTID -- 5.2 ARES -- 5.3 MONITOR -- 5.4 RUIAN -- 6 Related Work -- 7 Conclusions and Future Work -- References -- A Particle Swarm-Based Approach for Semantic Similarity Computation -- 1 Introduction -- 2 Background -- 2.1 Ontology-Based Similarity Measures -- 2.2 Ontology-Based Similarity Computation Scenarios -- 3 The Proposed Approach -- 3.1 Finding LCS as a Particle Swarm Optimization Problem -- 3.2 Fitness Scores of the LCS Evaluation -- 4 Experimental Evaluation -- 4.1 Benchmarks and Ontologies -- 4.2 Implementation -- 4.3 Experimental Results -- 5 Conclusion -- References -- Agent-Based Assistance in Ambient Assisted Living Through Reinforcement Learning and Semantic Technologies -- 1 Introduction -- 2 Related Work -- 3 Approach -- 3.1 Reinforcement Learning -- 3.2 Agent Semantics -- 3.3 Automated ``Semantification'' and Retrieval of Agent Policies -- 4 First Results and Conclusion -- References.
On the Need for Applications Aware Adaptive Middleware in Real-Time RDF Data Analysis (Short Paper) -- 1 Introduction -- 2 Formalization of RSP Features -- 3 Evaluation of RSP Engines Based on Formalized RSP Features -- 3.1 Latency and Breaking Points -- 3.2 Memory Consumption -- 4 Adaptive Middleware for RSP -- 4.1 Adaptive Layer -- 5 Conclusion -- References -- Learning Probabilistic Relational Models Using an Ontology of Transformation Processes -- 1 Introduction -- 2 Backgrounds -- 2.1 The Ontology of Transformation Processes -- 2.2 Probabilistic Relational Models -- 2.3 Learning PRMs -- 3 Learning a PRM Using an Ontology -- 3.1 Relational Schema Mapping -- 3.2 Our ON2PRM Algorithm -- 4 Experiments -- 4.1 Databases Generation -- 4.2 Results -- 5 Conclusion -- References -- ORDAIN: An Ontology for Trust Management in the Internet of Things -- Abstract -- 1 Introduction -- 2 Defining Trust and Reputation -- 3 ORDAIN Ontology -- 3.1 Ontology Contents -- 3.2 Ontology Implementation -- 4 Related Work -- 5 Conclusions -- Acknowledgments -- References -- APOPSIS: A Web-Based Platform for the Analysis of Structured Dialogues -- 1 Introduction -- 2 Background and Related Work -- 2.1 Computational Argumentation -- 2.2 Machine Learning -- 2.3 Existing Tools for Online Debates -- 3 A Methodology and Platform for Opinion Analysis -- 3.1 A Platform Methodology -- 3.2 Knowledge Map Representation -- 3.3 The MACE - Ontology Domain -- 3.4 Evaluating Opinions -- 4 Methodology: Analysing User-Generated Opinions -- 4.1 Clustering Analysis -- 5 Apopsis Implementation and Architecture Details -- 5.1 Debating Functionality -- 5.2 Debating on Different Levels of Discussions -- 5.3 Voting Mechanism -- 5.4 Searching Mechanism -- 5.5 General Features -- 5.6 Technical Information and Architecture -- 6 Conclusion and Future Work -- References.
Identifying Opinion Drivers on Social Media -- 1 Introduction -- 2 Related Literature -- 3 Driver Identification: Formal Model -- 4 Results and Evaluation -- 5 Conclusions -- References -- Representing Fashion Product Data with Schema.org: Approach and Use Cases -- Abstract -- 1 Introduction -- 2 Overview of Clothing-Related Information -- 2.1 Garment Sizes -- 2.2 Colors -- 2.3 Product Variants -- 2.4 Materials Information -- 2.5 Textile Care Recommendations --

2.6 Certifications -- 3 Clothing Product Information with Schema.org -- 3.1 Method -- 3.2 Representation of Basic Product Information -- 3.3 Representation of Clothing-Related Information -- 3.4 Representation of Textile Care Recommendations -- 3.5 Representation of Certifications -- 3.6 Other Information -- 3.7 Formalization and Results -- 4 Evaluation -- 5 Discussion and Conclusion -- 5.1 Related Work -- 5.2 Our Approach -- 5.3 Future Work -- References -- Semantic Modeling and Inference with Episodic Organization for Managing Personal Digital Traces -- 1 Introduction -- 2 Related Work -- 3 A Conceptual Model for Entities and Atomic Actions -- 4 Conceptual Model for scripts -- 5 Recognizing Script Instances from Documents -- 6 Summary -- References -- Linked Open Data for Linguists: Publishing the Hartmann von Aue-Portal in RDF -- Abstract -- 1 Introduction -- 1.1 Context -- 1.2 Problem Statement -- 1.3 Relevance -- 1.4 Organization of the Paper -- 2 Linked Open Data -- 3 The Hartmann von Aue-Portal -- 3.1 Purpose and Current Status -- 3.2 Data Model -- 3.3 Text Corpus Statistics -- 4 Approach -- 4.1 Challenges -- 4.2 Method -- 4.3 Implementation -- 4.4 Results -- 5 Evaluation -- 5.1 Formal Validation -- 5.2 Usage for Data-Driven Research -- 5.3 Interlinking of Hartmann von Aue Datasets with DBPedia -- 5.4 Reverse-Engineering -- 6 Discussion and Conclusion -- 6.1 Related Work.
6.2 Our Contribution.

Sommario/riassunto

This double volumes LNCS 10573-10574 constitutes the refereed proceedings of the Confederated International Conferences: Cooperative Information Systems, CoopIS 2017, Ontologies, Databases, and Applications of Semantics, ODBASE 2017, and Cloud and Trusted Computing, C&TC, held as part of OTM 2017 in October 2017 in Rhodes, Greece. The 61 full papers presented together with 19 short papers were carefully reviewed and selected from 180 submissions. The OTM program every year covers data and Web semantics, distributed objects, Web services, databases, information systems, enterprise workow and collaboration, ubiquity, interoperability, mobility, grid and high-performance computing.