

Resource Oriented Architecture Patterns For Webs Of Data Synthesis Lectures On The Semantic Web Theory And Technology

After a slow incubation period of nearly 15 years, a large and growing number of organizations now have one or more projects using the Semantic Web stack of technologies. The Web Ontology Language (OWL) is an essential ingredient in this stack, and the need for ontologists is increasing faster than the number and variety of available resources for learning OWL. This is especially true for the primary target audience for this book: modelers who want to build OWL ontologies for practical use in enterprise and government settings. The purpose of this book is to speed up the process of learning and mastering OWL. To that end, the focus is on the 30% of OWL that gets used 90% of the time. Others who may benefit from this book include technically oriented managers, semantic technology developers, undergraduate and post-graduate students, and finally, instructors looking for new ways to explain OWL. The book unfolds in a spiral manner, starting with the core ideas. Each subsequent cycle reinforces and expands on what has been learned in prior cycles and introduces new related ideas. Part 1 is a cook's tour of ontology and OWL, giving an informal overview of what things need to be said to build an ontology, followed by a detailed look at how to say them in OWL. This is illustrated using a healthcare example. Part 1 concludes with an explanation of some foundational ideas about meaning and semantics to prepare the reader for subsequent chapters. Part 2 goes into depth on properties and classes, which are the core of OWL. There are detailed descriptions of the main constructs that you are likely to need in every day modeling, including what inferences are sanctioned. Each is illustrated with real world examples. Part 3 explains and illustrates how to put OWL into practice, using examples in healthcare, collateral, and financial transactions. A small ontology is described for each, along with some key inferences. Key limitations of OWL are identified, along with possible workarounds. The final chapter gives a variety of practical tips and guidelines to send the reader on their way.

The World Wide Web is now deeply intertwined with our lives, and has become a catalyst for a data deluge, making vast amounts of data available online, at a click of a button. With Web 2.0, users are no longer passive consumers, but active publishers and curators of data. Hence, from science to food manufacturing, from data journalism to personal well-being, from social media to art, there is a strong interest in provenance, a description of what influenced an artifact, a data set, a document, a blog, or any resource on the Web and beyond. Provenance is a crucial piece of information that can help a consumer make a judgment as to whether something can be trusted. Provenance is no longer seen as a curiosity in art circles, but it is regarded as pragmatically, ethically, and methodologically crucial for our day-to-day data manipulation and curation activities on the Web. Following the recent publication of the PROV standard for provenance on the Web, which the two authors actively help shape in the Provenance Working Group at the World Wide Web Consortium, this Synthesis lecture is a hands-on introduction to PROV aimed at Web and linked data professionals. By means of recipes, illustrations, a website at www.provbook.org, and tools, it guides practitioners through a variety of issues related to provenance: how to generate provenance, publish it on the Web, make it discoverable, and how to utilize it. Equipped with this knowledge, practitioners will be in a position to develop novel applications that can bring open-ness, trust, and accountability. Table of Contents: Preface / Acknowledgments / Introduction / A Data Journalism Scenario / The PROV Ontology / Provenance Recipes / Validation, Compliance, Quality, Replay / Provenance Management / Conclusion / Bibliography / Authors' Biographies / Index

The practice of enterprise application development has benefited from the emergence of many new enabling technologies. Multi-tiered object-oriented platforms, such as Java and .NET, have become commonplace. These new tools and technologies are capable of building powerful applications, but they are not easily implemented. Common failures in enterprise applications often occur because their developers do not understand the architectural lessons that experienced object developers have learned. Patterns of Enterprise Application Architecture is written in direct response to the stiff challenges that face enterprise application developers. The author, noted object-oriented designer Martin Fowler, noticed that despite changes in technology--from Smalltalk to CORBA to Java to .NET--the same basic design ideas can be adapted and applied to solve common problems. With the help of an expert group of contributors, Martin distills over forty recurring solutions into patterns. The result is an indispensable handbook of solutions that are applicable to any enterprise application platform. This book is actually two books in one. The first section is a short tutorial on developing enterprise applications, which you can read from start to finish to understand the scope of the book's lessons. The next section, the bulk of the book, is a detailed reference to the patterns themselves. Each pattern provides usage and implementation information, as well as detailed code examples in Java or C#. The entire book is also richly illustrated with UML diagrams to further explain the concepts. Armed with this book, you will have the knowledge necessary to make important architectural decisions about building an enterprise application and the proven patterns for use when building them. The topics covered include · Dividing an enterprise application into layers · The major approaches to organizing business logic · An in-depth treatment of mapping between objects and relational databases · Using Model-View-Controller to organize a Web presentation · Handling concurrency for data that spans multiple transactions · Designing distributed object interfaces

In cooperation with experts and practitioners throughout the SOA community, best-selling author Thomas Erl brings together the de facto catalog of design patterns for SOA and service-orientation. More than three years in development and subjected to numerous industry reviews, the 85 patterns in this full-color book provide the most successful and proven design techniques to overcoming the most common and critical problems to achieving modern-day SOA. Through numerous examples, individually documented pattern profiles, and over 400 color illustrations, this book provides in-depth coverage of: • Patterns for the design, implementation, and governance of service inventories--collections of services representing individual service portfolios that can be independently modeled, designed, and evolved. • Patterns specific to service-level architecture which pertain to a wide range of design areas, including contract design, security, legacy encapsulation, reliability, scalability, and a variety of implementation and governance issues. • Service composition patterns that address the many aspects associated with combining services into aggregate distributed solutions, including topics such as runtime messaging and message design, inter-service security controls, and transformation. • Compound patterns (such as Enterprise Service Bus and Orchestration) and recommended pattern application sequences that establish foundational processes. The book begins by establishing SOA types that are referenced throughout the patterns and then form the basis of a final chapter that discusses the architectural impact of service-oriented computing in general. These chapters bookend the pattern catalog to provide a clear link between SOA design patterns, the strategic goals of service-oriented computing, different SOA types, and the service-orientation design paradigm. This book series is further supported by a series of resources sites, including soabooks.com, soaspecs.com, soapatterns.org, soamag.com, and soaposters.com.

The past ten years have seen a rapid growth in the numbers of people signing up to use Web-based social networks (hundreds of millions of new members are now joining the main services each year) with a large amount of content being shared on these networks (tens of billions of content items are shared each month). With this growth in usage and data being generated, there are many opportunities to discover the knowledge that is often inherent but somewhat hidden in these networks. Web mining techniques are being used to derive this hidden knowledge. In addition, the Semantic Web, including the Linked Data initiative to connect previously disconnected datasets, is making it possible to connect data from across various social spaces through common representations and agreed upon terms for people,

content items, etc. In this book, we detail some current research being carried out to semantically represent the implicit and explicit structures on the Social Web, along with the techniques being used to elicit relevant knowledge from these structures, and we present the mechanisms that can be used to intelligently mesh these semantic representations with intelligent knowledge discovery processes. We begin this book with an overview of the origins of the Web, and then show how web intelligence can be derived from a combination of web and Social Web mining. We give an overview of the Social and Semantic Webs, followed by a description of the combined Social Semantic Web (along with some of the possibilities it affords), and the various semantic representation formats for the data created in social networks and on social media sites. Provenance and provenance mining is an important aspect here, especially when data is combined from multiple services. We will expand on the subject of provenance and especially its importance in relation to social data. We will describe extensions to social semantic vocabularies specifically designed for community mining purposes (SIOCM). In the last three chapters, we describe how the combination of web intelligence and social semantic data can be used to derive knowledge from the Social Web, starting at the community level (macro), and then moving through group mining (meso) to user profile mining (micro). Table of Contents: Acknowledgments / Grant Aid / Introduction and the Web / Web Mining / The Social Web / The Semantic Web / The Social Semantic Web / Social Semantic Web Mining / Social Semantic Web Mining of Communities / Social Semantic Web Mining of Groups / Social Semantic Web Mining of Users / Conclusions / Bibliography / Authors' Biographies

The first volume of the POSA pattern series introduced a broad-spectrum of general-purpose patterns in software design and architecture. The second narrowed the focus to fundamental patterns for building sophisticated concurrent and networked software systems and applications. This volume uses design patterns to present techniques for implementing effective resource management in a system. The patterns are covered in detail making use of several examples providing directions to the readers on how to implement the presented patterns. Additionally, the volume presents a thorough introduction into resource management and a case study where the patterns are applied to the domain of mobile radio networks. The patterns are grouped by different areas of resource management and hence address the complete lifecycle of resources: resource acquisition, coordination and release.

Ontologies have become increasingly important as the use of knowledge graphs, machine learning, natural language processing (NLP), and the amount of data generated on a daily basis has exploded. As of 2014, 90% of the data in the digital universe was generated in the two years prior, and the volume of data was projected to grow from 3.2 zettabytes to 40 zettabytes in the next six years. The very real issues that government, research, and commercial organizations are facing in order to sift through this amount of information to support decision-making alone mandate increasing automation. Yet, the data profiling, NLP, and learning algorithms that are ground-zero for data integration, manipulation, and search provide less than satisfactory results unless they utilize terms with unambiguous semantics, such as those found in ontologies and well-formed rule sets. Ontologies can provide a rich "schema" for the knowledge graphs underlying these technologies as well as the terminological and semantic basis for dramatic improvements in results. Many ontology projects fail, however, due at least in part to a lack of discipline in the development process. This book, motivated by the Ontology 101 tutorial given for many years at what was originally the Semantic Technology Conference (SemTech) and then later from a semester-long university class, is designed to provide the foundations for ontology engineering. The book can serve as a course textbook or a primer for all those interested in ontologies.

The current drug development paradigm---sometimes expressed as, "One disease, one target, one drug"---is under question, as relatively few drugs have reached the market in the last two decades. Meanwhile, the research focus of drug discovery is being placed on the study of drug action on biological systems as a whole, rather than on individual components of such systems. The vast amount of biological information about genes and proteins and their modulation by small molecules is pushing drug discovery to its next critical steps, involving the integration of chemical knowledge with these biological databases. Systematic integration of these heterogeneous datasets and the provision of algorithms to mine the integrated datasets would enable investigation of the complex mechanisms of drug action; however, traditional approaches face challenges in the representation and integration of multi-scale datasets, and in the discovery of underlying knowledge in the integrated datasets. The Semantic Web, envisioned to enable machines to understand and respond to complex human requests and to retrieve relevant, yet distributed, data, has the potential to trigger system-level chemical-biological innovations. Chem2Bio2RDF is presented as an example of utilizing Semantic Web technologies to enable intelligent analyses for drug discovery. Table of Contents: Introduction / Data Representation and Integration Using RDF / Data Representation and Integration Using OWL / Finding Complex Biological Relationships in PubMed Articles using Bio-LDA / Integrated Semantic Approach for Systems Chemical Biology Knowledge Discovery / Semantic Link Association Prediction / Conclusions / References / Authors' Biographies

The Semantic Web is a young discipline, even if only in comparison to other areas of computer science. Nonetheless, it already exhibits an interesting history and evolution. This book is a reflection on this evolution, aiming to take a snapshot of where we are at this specific point in time, and also showing what might be the focus of future research. This book provides both a conceptual and practical view of this evolution, especially targeted at readers who are starting research in this area and as support material for their supervisors. From a conceptual point of view, it highlights and discusses key questions that have animated the research community: what does it mean to be a Semantic Web system and how is it different from other types of systems, such as knowledge systems or web-based information systems? From a more practical point of view, the core of the book introduces a simple conceptual framework which characterizes Intelligent Semantic Web Systems. We describe this framework, the components it includes, and give pointers to some of the approaches and technologies that might be used to implement them. We also look in detail at concrete systems falling under the category of Intelligent Semantic Web Systems, according to the proposed framework, allowing us to compare them, analyze their strengths and weaknesses, and identify the key fundamental challenges still open for researchers to tackle. Online social networks have already become a bridge connecting our physical daily life with the (web-based) information space. This connection produces a huge volume of data, not only about the information itself, but also about user behavior. The ubiquity of the social Web and the wealth of social data offer us unprecedented opportunities for studying the interaction patterns among users so as to understand the dynamic mechanisms underlying different networks, something that was previously difficult to explore due to the lack of available data. In this book, we present the architecture of the research for social network mining, from a microscopic point of view. We focus on investigating several key issues in social networks. Specifically, we begin with analytics of social interactions between users. The first kinds of questions we try to answer are: What are the fundamental factors that form the different categories of social ties? How have reciprocal relationships been developed from parasocial relationships? How do connected users further form groups? Another theme addressed in this book is the study of social influence. Social influence occurs when one's opinions, emotions, or behaviors are affected by others, intentionally or unintentionally. Considerable research has been conducted to verify the existence of social influence in various networks. However, few literature studies address how to quantify the strength of influence between users from different aspects. In Chapter 4 and in [138], we have studied how to model and predict user behaviors. One fundamental problem is distinguishing the effects of different social factors such as social influence, homophily, and individual's characteristics. We introduce a probabilistic model to address this problem. Finally, we use an academic social network, ArnetMiner, as an example to demonstrate how we apply the introduced technologies for mining real social networks. In this system, we try to mine knowledge from both the informative (publication) network and the social (collaboration) network, and to understand the interaction mechanisms between the two networks. The system has been in operation since 2006 and has already attracted millions of users from

more than 220 countries/regions.

Service-oriented architecture (SOA) is a flexible set of design principles used during the phases of systems development and integration in computing. A system based on a SOA will package functionality as a suite of interoperable services that can be used within multiple, separate systems from several business domains. SOA also generally provides a way for consumers of services, such as web-based applications, to be aware of available SOA-based services. For example, several disparate departments within a company may develop and deploy SOA services in different implementation languages; their respective clients will benefit from a well understood, well defined interface to access them. XML is commonly used for interfacing with SOA services, though this is not required. This book is your ultimate resource for Service-oriented architecture (SOA). Here you will find the most up-to-date information, analysis, background and everything you need to know. In easy to read chapters, with extensive references and links to get you to know all there is to know about Service-oriented architecture (SOA) right away, covering: Service-oriented architecture, Apache ActiveMQ, Apache Camel, Apache Qpid, Apache ServiceMix, Apache Synapse, Apache Tuscany, Apatar, Application Response Measurement, Boot image control, Comparison of business integration software, Canonical Model, Communications-enabled application, Corticon, Data element, DataNucleus, E-Biz Integrator, Enterprise application integration, Enterprise content management, Enterprise Integration Patterns, Enterprise messaging system, Enterprise service bus, EntireX, Event-driven architecture, Event-driven SOA, FuseSource Corp., Fuse ESB, Fuse Message Broker, GNU Enterprise, Governance Interoperability Framework, Guarana DSL, Hoox, IgniteXML, IGrafx, Information silo, Integrated software, Integration Objects, ISIS Papyrus, JBoss Enterprise SOA Platform, Jitterbit Integration Server, Loose coupling, Message-oriented middleware, Metaserver, Microsoft Enterprise Library, Openadaptor, OpenBRR, OpenGate, Oracle Enterprise Service Bus, Oracle SOA Suite, Orchestration (computing), Pervasive business intelligence, Petals ESB, Sarvega, Search-based application, SEEBURGER, Semantic service-oriented architecture, Semantic translation, SEMCI, Service Component Architecture, Service Oriented Architecture Fundamentals, Service-oriented architecture implementation framework, SOALIB, Talend, ThoughtWorks, Total cost of ownership, Tryton, Universal integration platform, Verastream Host Integrator, Virtuoso Universal Server, Web-oriented architecture, WebORB Integration Server, While You Were Out (Cloud application), WS-CAF, Adaptive Services Grid, Application fabrication, B2B Gateway, Barracuda Networks, Boomerang Software Framework, Business Process Network, Canonical Protocol Pattern, Canonical Schema pattern, Composite application, Denodo, Differentiated service, Digital Nervous System, Domain Inventory Pattern, Enterprise Inventory, Enterprise Service Layer, Entity Abstraction Pattern, Event-Driven Messaging, Experticity, Freightgate, Intel SOA Products Division, JackBe, Logic Centralization Pattern, Machine-to-Machine, Midas Kapiti, Multitenancy, Mushroom Networks, MVaaS, Net-Centric Enterprise Services, Network-Centric Service-Oriented Enterprise (NCSE), OASIS SOA Reference Model, Open Knowledge Initiative, Open Mashup Alliance, Open Service Interface Definitions, Opti-Time Company, Oslo (Microsoft), Postini, Reliable messaging, S-RAMP, SAP Enterprise Architecture Framework, Service (systems architecture), Service Abstraction, Service Autonomy Principle...and much more This book explains in-depth the real drivers and workings of Service-oriented architecture (SOA). It reduces the risk of your technology, time and resources investment decisions by enabling you to compare your understanding of Service-oriented architecture (SOA) with the objectivity of experienced professionals.

"A stereotype of computer science textbooks is that they are dry, boring, and sometimes even intimidating. As a result, they turn students' interests off from the subject matter instead of enticing them into it. This textbook is the opposite of such a stereotype. The author presents the subject matter in a refreshing story-telling style and aims to bring the Internet-generation of students closer to her stories." --Yingcai Xiao, The University of Akron Introduction to Middleware: Web Services, Object Components, and Cloud Computing provides a comparison of different middleware technologies and the overarching middleware concepts they are based on. The various major paradigms of middleware are introduced and their pros and cons are discussed. This includes modern cloud interfaces, including the utility of Service Oriented Architectures. The text discusses pros and cons of RESTful vs. non-RESTful web services, and also compares these to older but still heavily used distributed object/component middleware. The text guides readers to select an appropriate middleware technology to use for any given task, and to learn new middleware technologies as they appear over time without being greatly overwhelmed by any new concept. The book begins with an introduction to different distributed computing paradigms, and a review of the different kinds of architectures, architectural styles/patterns, and properties that various researchers have used in the past to examine distributed applications and determine the quality of distributed applications. Then it includes appropriate background material in networking and the web, security, and encoding necessary to understand detailed discussion in this area. The major middleware paradigms are compared, and a comparison methodology is developed. Readers will learn how to select a paradigm and technology for a particular task, after reading this text. Detailed middleware technology review sections allow students or industry practitioners working to expand their knowledge to achieve practical skills based on real projects so as to become well-functional in that technology in industry. Major technologies examined include: RESTful web services (RESTful cloud interfaces such as OpenStack, AWS EC2 interface, CloudStack; AJAX, JAX-RS, ASP.NET MVC and ASP.NET Core), non-RESTful (SOAP and WSDL-based) web services (JAX-WS, Windows Communication Foundation), distributed objects/ components (Enterprise Java Beans, .NET Remoting, CORBA). The book presents two projects that can be used to illustrate the practical use of middleware, and provides implementations of these projects over different technologies. This versatile and class-tested textbook is suitable (depending on chapters selected) for undergraduate or first-year graduate courses on client server architectures, middleware, and cloud computing, web services, and web programming.

The book aims to provide a broad overview of various topics of the Internet of Things (IoT) from the research and development priorities to enabling technologies, architecture, security, privacy, interoperability and industrial applications. It is intended to be a stand-alone book in a series that covers the Internet of Things activities of the IERC - Internet of Things European Research Cluster - from technology to international cooperation and the global "state of play." The book builds on the ideas put forward by the European Research Cluster on the Internet of Things Strategic Research and Innovation Agenda and presents views and state of the art results on the challenges facing the research, development and deployment of IoT at the global level. Today we see the integration of Industrial, Business and Consumer Internet which is bringing together the Internet of People, Internet of Things, Internet of Energy, Internet of Vehicles, Internet of Media, Services and Enterprises in forming the backbone of the digital economy, the digital society and the foundation for the future knowledge and innovation based economy. These developments are supporting solutions for the emerging challenges of public health, aging population, environmental protection and climate change, the conservation of energy and scarce materials, enhancements to safety and security and the continuation and growth of economic prosperity. Penetration of smartphones and advances in nanoelectronics, cyber-physical systems, wireless communication, software, and Cloud computing technology will be the main drivers for IoT development. The IoT contribution is seen in the increased value of information created by the number of interconnections among things and the transformation of the processed information into knowledge shared into the Internet of Everything. The connected devices are part of ecosystems connecting people, processes, data, and things which are communicating in the Cloud using the increased storage and computing power while attempting to standardize communication and metadata. In this context, the next generation of Cloud computing technologies will need to be flexible enough to scale autonomously, adaptive enough to handle constantly changing connections and resilient enough to stand up to the huge flows of data that will occur. In 2025, analysts forecast that there will be six devices per human on the planet, which means around 50 billion more connected devices over the next 12 years. The Internet of Things market is connected to this anticipated device growth from industrial Machine to Machine (M2M) systems,

smart meters and wireless sensors. Internet of Things technology will generate new services and new interfaces by creating smart environments and smart spaces with applications ranging from Smart Cities, Smart Transport, Buildings, Energy, Grid, to Smart Health and Life.

In recent years, several knowledge bases have been built to enable large-scale knowledge sharing, but also an entity-centric Web search, mixing both structured data and text querying. These knowledge bases offer machine-readable descriptions of real-world entities, e.g., persons, places, published on the Web as Linked Data. However, due to the different information extraction tools and curation policies employed by knowledge bases, multiple, complementary and sometimes conflicting descriptions of the same real-world entities may be provided. Entity resolution aims to identify different descriptions that refer to the same entity appearing either within or across knowledge bases. The objective of this book is to present the new entity resolution challenges stemming from the openness of the Web of data in describing entities by an unbounded number of knowledge bases, the semantic and structural diversity of the descriptions provided across domains even for the same real-world entities, as well as the autonomy of knowledge bases in terms of adopted processes for creating and curating entity descriptions. The scale, diversity, and graph structuring of entity descriptions in the Web of data essentially challenge how two descriptions can be effectively compared for similarity, but also how resolution algorithms can efficiently avoid examining pairwise all descriptions. The book covers a wide spectrum of entity resolution issues at the Web scale, including basic concepts and data structures, main resolution tasks and workflows, as well as state-of-the-art algorithmic techniques and experimental trade-offs. This proceedings volume contains selected revised and extended research articles written by researchers who participated in the World Congress on Engineering and Computer Science 2015, held in San Francisco, USA, 21-23 October 2015. Topics covered include engineering mathematics, electrical engineering, circuits, communications systems, computer science, chemical engineering, systems engineering, manufacturing engineering, and industrial applications. The book offers the reader an overview of the state of the art in engineering technologies, computer science, systems engineering and applications, and will serve as an excellent reference work for researchers and graduate students working in these fields.

As more and more devices become interconnected through the Internet of Things (IoT), there is an even greater need for this book, which explains the technology, the internetworking, and applications that are making IoT an everyday reality. The book begins with a discussion of IoT "ecosystems" and the technology that enables them, which includes: Wireless Infrastructure and Service Discovery Protocols Integration Technologies and Tools Application and Analytics Enablement Platforms A chapter on next-generation cloud infrastructure explains hosting IoT platforms and applications. A chapter on data analytics throws light on IoT data collection, storage, translation, real-time processing, mining, and analysis, all of which can yield actionable insights from the data collected by IoT applications. There is also a chapter on edge/fog computing. The second half of the book presents various IoT ecosystem use cases. One chapter discusses smart airports and highlights the role of IoT integration. It explains how mobile devices, mobile technology, wearables, RFID sensors, and beacons work together as the core technologies of a smart airport. Integrating these components into the airport ecosystem is examined in detail, and use cases and real-life examples illustrate this IoT ecosystem in operation. Another in-depth look is on envisioning smart healthcare systems in a connected world. This chapter focuses on the requirements, promising applications, and roles of cloud computing and data analytics. The book also examines smart homes, smart cities, and smart governments. The book concludes with a chapter on IoT security and privacy. This chapter examines the emerging security and privacy requirements of IoT environments. The security issues and an assortment of surmounting techniques and best practices are also discussed in this chapter.

RDF and Linked Data have broad applicability across many fields, from aircraft manufacturing to zoology. Requirements for detecting bad data differ across communities, fields, and tasks, but nearly all involve some form of data validation. This book introduces data validation and describes its practical use in day-to-day data exchange. The Semantic Web offers a bold, new take on how to organize, distribute, index, and share data. Using Web addresses (URIs) as identifiers for data elements enables the construction of distributed databases on a global scale. Like the Web, the Semantic Web is heralded as an information revolution, and also like the Web, it is encumbered by data quality issues. The quality of Semantic Web data is compromised by the lack of resources for data curation, for maintenance, and for developing globally applicable data models. At the enterprise scale, these problems have conventional solutions. Master data management provides an enterprise-wide vocabulary, while constraint languages capture and enforce data structures. Filling a need long recognized by Semantic Web users, shapes languages provide models and vocabularies for expressing such structural constraints. This book describes two technologies for RDF validation: Shape Expressions (ShEx) and Shapes Constraint Language (SHACL), the rationales for their designs, a comparison of the two, and some example applications.

Linked Data (LD) is a well-established standard for publishing and managing structured information on the Web, gathering and bridging together knowledge from different scientific and commercial domains. The development of Linked Data Visualization techniques and tools has been adopted as the established practice for the analysis of this vast amount of information by data scientists, domain experts, business users, and citizens. This book covers a wide spectrum of visualization topics, providing an overview of the recent advances in this area, focusing on techniques, tools, and use cases of visualization and visual analysis of LD. It presents core concepts related to data visualization and LD technologies, techniques employed for data visualization based on the characteristics of data, techniques for Big Data visualization, tools and use cases in the LD context, and, finally, a thorough assessment of the usability of these tools under different scenarios. The purpose of this book is to offer a complete guide to the evolution of LD visualization for interested readers from any background and to empower them to get started with the visual analysis of such data. This book can serve as a course textbook or as a primer for all those interested in LD and data visualization.

A complete practitioner's catalog of proven domain services design solutions that can help any organization leverage SOA's full benefits * *Provides a vocabulary of proven SOA design solutions, with concrete examples and code that is easy for architects to adapt and implement. *By Rob Daigneau, one of the industry's leading experts in complex systems integration. *Helps architects and IT leaders accurately set stakeholder expectations for major SOA initiatives. Service-oriented architectures are typically called upon to deliver two general categories of services: enterprise services and domain services. Enterprise services are essentially composite services that typically leverage technologies such as message-oriented middleware. Domain services are the building blocks these composites depend upon. Each service category is best served by a distinct set of design

solutions. This is the first book to systematically identify and explain best practice patterns for domain services. Rob Daigneau expands upon the Service Layer concept (covered expertly by Fowler in *Patterns of Enterprise Application Architecture*) domain services can be used with *Enterprise Integration Patterns* (made famous by Hohpe and Woolf). Daigneau begins by reviewing SOA concepts, illuminating the distinctions between enterprise and domain services, and identifying key relationships between domain services and other pattern groups. Next, he introduces each essential pattern for creating and delivering domain services, providing a vocabulary of design solutions that architects and other IT professionals can implement by referencing and adapting the concrete examples he supplies.

Enterprise Architecture (EA) is typically an aggregate of the business, application, data, and infrastructure architectures of any forward-looking enterprise. Due to constant changes and rising complexities in the business and technology landscapes, producing sophisticated architectures is on the rise. Architectural patterns are gaining a lot ...

As data management and integration continue to evolve rapidly, storing all your data in one place, such as a data warehouse, is no longer scalable. In the very near future, data will need to be distributed and available for several technological solutions. With this practical book, you'll learn how to migrate your enterprise from a complex and tightly coupled data landscape to a more flexible architecture ready for the modern world of data consumption. Executives, data architects, analytics teams, and compliance and governance staff will learn how to build a modern scalable data landscape using the Scaled Architecture, which you can introduce incrementally without a large upfront investment. Author Piethein Strengholt provides blueprints, principles, observations, best practices, and patterns to get you up to speed. Examine data management trends, including technological developments, regulatory requirements, and privacy concerns Go deep into the Scaled Architecture and learn how the pieces fit together Explore data governance and data security, master data management, self-service data marketplaces, and the importance of metadata

"Every developer working with the Web needs to read this book." -- David Heinemeier Hansson, creator of the Rails framework "RESTful Web Services finally provides a practical roadmap for constructing services that embrace the Web, instead of trying to route around it." -- Adam Trachtenberg, PHP author and EBay Web Services Evangelist You've built web sites that can be used by humans. But can you also build web sites that are usable by machines? That's where the future lies, and that's what RESTful Web Services shows you how to do. The World Wide Web is the most popular distributed application in history, and Web services and mashups have turned it into a powerful distributed computing platform. But today's web service technologies have lost sight of the simplicity that made the Web successful. They don't work like the Web, and they're missing out on its advantages. This book puts the "Web" back into web services. It shows how you can connect to the programmable web with the technologies you already use every day. The key is REST, the architectural style that drives the Web. This book: Emphasizes the power of basic Web technologies -- the HTTP application protocol, the URI naming standard, and the XML markup language Introduces the Resource-Oriented Architecture (ROA), a common-sense set of rules for designing RESTful web services Shows how a RESTful design is simpler, more versatile, and more scalable than a design based on Remote Procedure Calls (RPC) Includes real-world examples of RESTful web services, like Amazon's Simple Storage Service and the Atom Publishing Protocol Discusses web service clients for popular programming languages Shows how to implement RESTful services in three popular frameworks -- Ruby on Rails, Restlet (for Java), and Django (for Python) Focuses on practical issues: how to design and implement RESTful web services and clients This is the first book that applies the REST design philosophy to real web services. It sets down the best practices you need to make your design a success, and the techniques you need to turn your design into working code. You can harness the power of the Web for programmable applications: you just have to work with the Web instead of against it. This book shows you how.

REST architecture (style) is a pivot of distributed systems, simplify data integration amongst modern and legacy applications leverages through the RESTful paradigm. This book is fully loaded with many RESTful API patterns, samples, hands-on implementations and also discuss the capabilities of many REST API frameworks for Java, Scala, Python and Go

Web browsing would not be what it is today without the use of Service-Oriented Architecture (SOA). Although much has been written about SOA methodology, this emerging platform is continuously under development. *Exploring Enterprise Service Bus in the Service-Oriented Architecture Paradigm* is a detailed reference source that examines current aspects and research methodologies that enable enterprise service bus to unify and connect services efficiently on a common platform. Featuring relevant topics such as SOA reference architecture, grid computing applications, complex event computing, and java business integration, this is an ideal resource for all practitioners, academicians, graduate students, and researchers interested in the discoveries on the relationship that Service-Oriented architecture and enterprise service bus share.

In this truly unique technical book, today's leading software architects present valuable principles on key development issues that go way beyond technology. More than four dozen architects -- including Neal Ford, Michael Nygard, and Bill de hOra -- offer advice for communicating with stakeholders, eliminating complexity, empowering developers, and many more practical lessons they've learned from years of experience. Among the 97 principles in this book, you'll find useful advice such as: Don't Put Your Resume Ahead of the Requirements (Nitin Borwankar) Chances Are, Your Biggest Problem Isn't Technical (Mark Ramm) Communication Is King; Clarity and Leadership, Its Humble Servants (Mark Richards) Simplicity Before Generality, Use Before Reuse (Kevlin Henney) For the End User, the Interface Is the System (Vinayak Hegde) It's Never Too Early to Think About Performance (Rebecca Parsons) To be successful as a software architect, you need to master both business and technology. This book tells you what top software architects think is important and how they approach a project. If you want to enhance your career, *97 Things Every Software Architect Should Know* is essential reading.

This book introduces core natural language processing (NLP) technologies to non-experts in an easily accessible way, as a series of building blocks that lead the user to understand key technologies, why they are required, and how to integrate them into Semantic Web applications. Natural language processing and Semantic Web technologies have different, but complementary roles in data management. Combining these two technologies enables structured and unstructured data to merge seamlessly. Semantic Web technologies aim to convert unstructured data to meaningful representations, which benefit enormously from the use of NLP technologies, thereby enabling applications such as connecting text to Linked Open Data, connecting texts to each other, semantic searching, information visualization, and modeling of user behavior in online networks. The first half of this book describes the basic NLP processing tools: tokenization, part-of-speech tagging, and morphological analysis, in addition to the main tools required for an information extraction system (named entity recognition and relation extraction) which build on these components. The second half of the book explains how Semantic Web and NLP technologies can enhance each other, for example via semantic annotation, ontology linking, and population. These chapters also discuss sentiment analysis, a key component in making sense of textual data, and the difficulties of performing NLP on social media, as well as some proposed solutions. The book finishes by investigating some applications of these tools, focusing on semantic search and visualization, modeling user behavior, and an outlook on the future.

Build straightforward and maintainable APIs to create services that are usable and maintainable. Although this book focuses on distributed services, it also emphasizes how the core principles apply even to pure OOD and OOP constructs. The overall context of Creating Maintainable APIs is to classify the topics into four main areas: classes and interfaces, HTTP REST APIs, messaging APIs, and message payloads (XML, JSON and JSON API as well as Apache Avro). What You Will Learn Use object-oriented design constructs and their APIs Create and manage HTTP REST APIs Build and manage maintainable messaging APIs, including the use of Apache Kafka as a principal messaging hub Handle message payloads via JSON Who This Book Is For Any level software engineers and very experienced programmers.

This book explains how to combine and exploit sensor networks and internet-of-things (IoT) technologies and Web-service design patterns to enrich and integrate Building Information Models (BIMs). It provides approaches and software architectures for facilitating the interaction with (and between) BIMs through Web services, and for enabling and facilitating the fusion of the information residing in such models or of information acquired from IoT technologies. The proposed software architectures are presented in the form of design patterns. This information fusion will facilitate many novel application fields ranging from emergency response, to urban monitoring and surveillance, and to smart buildings. The book consists of 8 chapters. The first 2 chapters focus on the basics of BIMs, while chapter 3 presents fundamental service-oriented architecture patterns for complex information models. Subsequently, chapters 4 and 5 elaborate on the hardware and software side of IoT, with a special focus on their use for BIMs. Chapter 6 provides advanced SOA patterns for BIMs, while chapter 7 details patterns for IoT, and for BIM and IoT information fusion. Lastly, chapter 8 summarizes the work and provides an outlook on promising future developments. Overall, the book will be beneficial for researchers and developers in the fields of building information models, IoT applications, and systems integration.

The surge of interest in the REpresentational State Transfer (REST) architectural style, the Semantic Web, and Linked Data has resulted in the development of innovative, flexible, and powerful systems that embrace one or more of these compatible technologies. However, most developers, architects, Information Technology managers, and platform owners have only been exposed to the basics of resource-oriented architectures. This book is an attempt to catalog and elucidate several reusable solutions that have been seen in the wild in the now increasingly familiar "patterns book" style. These are not turn key implementations, but rather, useful strategies for solving certain problems in the development of modern, resource-oriented systems, both on the public Web and within an organization's firewalls.

Modern web applications deserve modern tools. Harness the JVM's rich infrastructure while taking advantage of the expressive power and brisk performance of a modern functional language. Exploit Clojure's unique advantages for web development. Step by step, apply the fundamentals of programming in Clojure to build real-world, professional web applications. This edition features new libraries, tools, and best practices, and focuses on developing modern single-page applications. Stop developing web apps with yesterday's tools. Today, developers are increasingly adopting Clojure as a web-development platform. See for yourself what makes Clojure so desirable, as you create a series of web apps of growing complexity, exhibiting the full process of web development using a modern functional language. Journey through all the steps in developing a rich Picture Gallery web application--from conception to packaging and deployment. You'll work hands-on with Clojure and build real-world, professional web apps. This fully updated second edition reveals the changes in the rapidly evolving Clojure ecosystem. Get up to speed on the many new libraries, tools, and best practices. Gain expertise in the popular Ring/Compojure stack using the Luminus framework. Learn how Clojure works with databases and speeds development of RESTful services. See why ClojureScript is rapidly becoming a popular front-end platform, and use ClojureScript with the popular Reagent library to build single-page applications. This book is for you, whether you're already familiar with Clojure or if you're completely new to the language. What You Need: The latest JVM, Clojure 1.6+, and the Leiningen build tool, as well as an editor such as Emacs, IntelliJ, Eclipse, Light Table, or VI.

This book is a guide to designing and building knowledge graphs from enterprise relational databases in practice. It presents a principled framework centered on mapping patterns to connect relational databases with knowledge graphs, the roles within an organization responsible for the knowledge graph, and the process that combines data and people. The content of this book is applicable to knowledge graphs being built either with property graph or RDF graph technologies. Knowledge graphs are fulfilling the vision of creating intelligent systems that integrate knowledge and data at large scale. Tech giants have adopted knowledge graphs for the foundation of next-generation enterprise data and metadata management, search, recommendation, analytics, intelligent agents, and more. We are now observing an increasing number of enterprises that seek to adopt knowledge graphs to develop a competitive edge. In order for enterprises to design and build knowledge graphs, they need to understand the critical data stored in relational databases. How can enterprises successfully adopt knowledge graphs to integrate data and knowledge, without boiling the ocean? This book provides the answers.

This timely volume provides a review of the state-of-the-art frameworks and methodologies for connecting diverse objects and devices according to the vision for an Internet of Things (IoT). A

specific focus is placed on the communication, security, and privacy aspects of device connectivity in distributed environments. Insights and case studies are provided by an authoritative selection of contributors of international repute into the latest research advances and practical approaches with respect to the connectivity of heterogeneous smart and sensory devices. Topics and features: Examines aspects of device connectivity within the IoT Presents a resource-based architecture for IoT, and proposes a resource management framework for corporate device clouds Reviews integration approaches for the IoT environment, and discusses performance optimization of intelligent home networks Introduces a novel solution for interoperable data management in multi-clouds, and suggests an approach that addresses the debate over network neutrality in the IoT Describes issues of data security, privacy, access control, and authentication in the distributed IoT environment Reviews the evolution of VANETs in relation to the Internet of Vehicles, and provides a perspective on developing smart sustainable cities This invaluable text/reference will be of great benefit to a broad audience, from students and researchers interested in the IoT vision, to practicing communication engineers and network security specialists.

Business and IT organizations are currently embracing new strategically sound concepts in order to be more customer-centric, competitive, and cognitive in their daily operations. While useful, the various software tools, pioneering technologies, as well as their unique contributions largely go unused due to the lack of information provided on their special characteristics. Novel Practices and Trends in Grid and Cloud Computing is a collection of innovative research on the key concerns of cloud computing and how they are being addressed, as well as the various technologies and tools empowering cloud theory to be participative, penetrative, pervasive, and persuasive. While highlighting topics including cyber security, smart technology, and artificial intelligence, this book is ideally designed for students, researchers, and business managers on the lookout for innovative IT solutions for all the business automation software and improvisations of computational technologies.

The dramatic progress of smartphone technologies has ushered in a new era of mobile sensing, where traditional wearable on-body sensors are being rapidly superseded by various embedded sensors in our smartphones. For example, a typical smartphone today, has at the very least a GPS, WiFi, Bluetooth, triaxial accelerometer, and gyroscope. Alongside, new accessories are emerging such as proximity, magnetometer, barometer, temperature, and pressure sensors. Even the default microphone can act as an acoustic sensor to track noise exposure for example. These sensors act as a "lens" to understand the user's context along different dimensions. Data can be passively collected from these sensors without interrupting the user. As a result, this new era of mobile sensing has fueled significant interest in understanding what can be extracted from such sensor data both instantaneously as well as considering volumes of time series from these sensors. For example, GPS logs can be used to determine automatically the significant places associated to a user's life (e.g., home, office, shopping areas). The logs may also reveal travel patterns, and how a user moves from one place to another (e.g., driving or using public transport). These may be used to proactively inform the user about delays, relevant promotions from shops, in his "regular" route. Similarly, accelerometer logs can be used to measure a user's average walking speed, compute step counts, gait identification, and estimate calories burnt per day. The key objective is to provide better services to end users. The objective of this book is to inform the reader of the methodologies and techniques for extracting meaningful information (called "semantics") from sensors on our smartphones. These techniques form the cornerstone of several application areas utilizing smartphone sensor data. We discuss technical challenges and algorithmic solutions for modeling and mining knowledge from smartphone-resident sensor data streams. This book devotes two chapters to dive deep into a set of highly available, commoditized sensors---the positioning sensor (GPS) and motion sensor (accelerometer). Furthermore, this book has a chapter devoted to energy-efficient computation of semantics, as battery life is a major concern on user experience.

This book describes a set of methods, architectures, and tools to extend the data pipeline at the disposal of developers when they need to publish and consume data from Knowledge Graphs (graph-structured knowledge bases that describe the entities and relations within a domain in a semantically meaningful way) using SPARQL, Web APIs, and JSON. To do so, it focuses on the paradigmatic cases of two middleware software packages, grlc and SPARQL Transformer, which automatically build and run SPARQL-based REST APIs and allow the specification of JSON schema results, respectively. The authors highlight the underlying principles behind these technologies—query management, declarative languages, new levels of indirection, abstraction layers, and separation of concerns—, explain their practical usage, and describe their penetration in research projects and industry. The book, therefore, serves a double purpose: to provide a sound and technical description of tools and methods at the disposal of publishers and developers to quickly deploy and consume Web Data APIs on top of Knowledge Graphs; and to propose an extensible and heterogeneous Knowledge Graph access infrastructure that accommodates a growing ecosystem of querying paradigms.

WebAssembly: The Definitive Guide is a thorough and accessible introduction to one of the most transformative technologies hitting our industry. What started as a way to use languages other than JavaScript in the browser has evolved into a comprehensive path toward portability, performance, increased security, and greater code reuse across an impressive collection of deployment targets. Author Brian Sletten introduces elements of this technology incrementally while building to several concrete, code-driven examples of practical, cutting-edge WebAssembly uses. Whether you work with enterprise software or embedded systems, or in entertainment, scientific computing, or startup environments, you'll learn how WebAssembly can have a positive impact on the way you develop software. Use WebAssembly to increase code portability across platforms Reuse more of your software assets in a wider number of deployment targets Learn how WebAssembly increases protection against prominent security attacks Use WebAssembly to deploy legacy code in web environments Increase your user base across languages and development environments Integrate JavaScript code with other languages and environments to improve performance, security, and productivity Learn how WebAssembly will affect your career as software developer

This book describes OCLC's contributions to the transformation of the Internet from a web of documents to a Web of Data. The new Web is a growing 'cloud' of interconnected resources that identify the things people want to know about when they approach the Internet with an information need. The linked data architecture has achieved critical mass just as it has become clear that library standards for resource description are nearing obsolescence. Working for the world's largest library cooperative, OCLC researchers have been active participants in the development of next generation standards for library resource description. By engaging with an international community of library and Web standards experts, they have published some of the most widely used RDF datasets representing library collections and librarianship. This book focuses on the conceptual and technical challenges involved in publishing linked data derived from traditional library metadata. This transformation is a high priority because most searches for information start not in the library, nor even in a Web-accessible library catalog, but elsewhere on the Internet. Modeling data in a form that the broader Web understands will project the value of libraries into the Digital Information Age. The exposition is aimed at librarians, archivists, computer scientists, and other professionals interested in modeling bibliographic descriptions as linked data. It aims to achieve a balanced treatment of theory, technical detail, and practical application.

This book provides a comprehensive and accessible introduction to knowledge graphs, which have recently garnered notable attention from both industry and academia. Knowledge graphs are founded on the principle of applying a graph-based abstraction to data, and are now broadly deployed in scenarios that require integrating and extracting value from multiple, diverse sources of data at large scale. The

book defines knowledge graphs and provides a high-level overview of how they are used. It presents and contrasts popular graph models that are commonly used to represent data as graphs, and the languages by which they can be queried before describing how the resulting data graph can be enhanced with notions of schema, identity, and context. The book discusses how ontologies and rules can be used to encode knowledge as well as how inductive techniques—based on statistics, graph analytics, machine learning, etc.—can be used to encode and extract knowledge. It covers techniques for the creation, enrichment, assessment, and refinement of knowledge graphs and surveys recent open and enterprise knowledge graphs and the industries or applications within which they have been most widely adopted. The book closes by discussing the current limitations and future directions along which knowledge graphs are likely to evolve. This book is aimed at students, researchers, and practitioners who wish to learn more about knowledge graphs and how they facilitate extracting value from diverse data at large scale. To make the book accessible for newcomers, running examples and graphical notation are used throughout. Formal definitions and extensive references are also provided for those who opt to delve more deeply into specific topics.

This book fills a gap between high-level overview texts that are often too general and low-level detail oriented technical handbooks that lose sight the "big picture". This book discusses SOA from the low-level perspective of middleware, various XML-based technologies, and basic service design. It also examines broader implications of SOA, particularly where it intersects with business process management and process modeling. Concrete overviews will be provided of the methodologies in those fields, so that students will have a hands-on grasp of how they may be used in the context of SOA.

Resource-Oriented Architecture Patterns for Webs of Data Morgan & Claypool Publishers

[Copyright: d4090eb0b4fda029f23c234734465d60](#)