

As applications have evolved to be more robust and complex, so has their data. Applications have almost entirely become solely dependent on their data. Data powers our applications and has helped them evolve to be the feature-full products we've come to love. In addition to the more complex requirements of software, developers have also begun shifting their approach for creating applications from the existing waterfall model. Historically, the problem of data has been solved through the use of one of the many flavors of relational databases such as MySQL, PostgreSQL, and Oracle. But, given how our applications and their data is evolving, the existing relational databases may no longer be the "one size fits all" solution that they have been before. A new breed of databases not based on the existing relational model have started to catch the attention of developers. These databases have been categorized as Not Only SQL databases or NoSQL databases. NoSQL databases include distributed key-value stores such as Redis and Riak, document-oriented databases such as MongoDB and CouchDB, and graph databases such as Neo4j. As the requirements of our applications have shifted, from the amount of data stored, agile software development, and requirements for availability, NoSQL databases have gradually become the more popular solution for storing data. Each type of database has proven to be beneficial for a set of use cases. Distributed key-value stores such as Redis have become popular to be used as an intermediate high availability and performance data stores. Document-oriented databases such as MongoDB have also become popular due to their rich feature-set, ability to shard data, and schema-less nature which allows developers to focus on applications without worrying about schema. Unfortunately, outside of their use cases, document and distributed key-value store databases, haven't helped in the ability to solve complex relationships. A use case exists for databases that easily allow developers to model complex relationships amongst data. While the original solution to this problem may have been relational databases, one may argue that relational databases aren't entirely relational and do not lend themselves to modeling complex relationships. While relational databases have partially solved this problem, a complete solution now exists with graph databases such as Neo4J. Graph databases combine the schema-less nature of other NoSQL databases but relationships are first class citizens and can easily be created to model real-world data. Given this well defined use case and a glaring problem that needs to be solved, this thesis will explore the use of a graph database for a non-trivial application. The application that will be examined is the Electronic Assessment System (EAS) developed by the Pioneering Technology Group (PTG) at California State University Northridge (CSUN). EAS is one of the premier applications developed by PTG which will be used by CSUN professors, students, and staff to assess student learning through their higher educational careers while attending CSUN. This thesis will explore, for the same set of data containing complex relationships, how would data be modelled in a graph database, how do the queries between each database compare to one another, and how it would have helped aid development within an Agile environment.

Understand the language and vocabulary of Data Architecture. The Data Architecture field is rife with terms that have become "fashionable". Some of the terms began with very specific, specialized, meanings – but as their use spread, they lost the precision of their technical definitions and become, well, "buzzwords". A buzzword is "a word or expression from a particular subject area that has become fashionable because it has been used a lot". Compliance is "the obeying of an accepted principle or instruction that states the way things are or should be done." The assignment is to take buzzwords and follow rules to use them correctly. We cut through the hype to arrive at buzzword compliance – the state where you fully understand the words that in fact have real meaning in the data architecture industry. This book will rationalize the various ways all these terms are defined. Of necessity, the book must address all aspects of describing an enterprise and its data management technologies. This includes a wide range of subjects, from entity/relationship modeling, through the semantic web, to database issues like relational and "beyond relational" ("NoSQL") approaches. In each case, the definitions for the subject are meant to be detailed enough to make it possible to understand basic principles—while recognizing that a full understanding will require consulting the sources where they are more completely described. The book's Glossary contains a catalogue of definitions and its Bibliography contains a comprehensive set of references.

"Graph is a commonly used data structure for modeling complex data such as chemical molecules, images, social networks, and XML documents. This complex data is stored using a set of graphs, known as graph database D. To speed up query answering on graph databases, indexes are commonly used. State-of-the-art graph database indexes do not adapt or scale well to dynamic graph database use; they are static, and their ability to prune possible search responses to meet user needs worsens over time as databases change and grow. Users can re-mine indexes to gain some improvement, but it is time consuming. Users must also tune numerous parameters on an ongoing basis to optimize performance and can inadvertently worsen the query response time if they do not choose parameters wisely. Recently, a one-pass algorithm has been developed to enhance the performance of these indexes in part by using the algorithm to update them regularly. However, there are some drawbacks, most notably the need to make updates as the query workload changes. We propose a new index based on graph-coarsening to speed up query answering time in dynamic graph databases. Our index is parameter-free, query-independent, scalable, small enough to store in the main memory, and is simpler and less costly to maintain for database updates. We conducted an extensive sets of experiments on two types of databases, i.e., chemical and social network databases, to compare our graph-coarsening based index vs. hybrid-indexes as follows. First, we considered no database updates or query workload changes (static graph databases) and compared the indexes according to query answering time and index size for different minSup values. Second, we compared the indexes in the case of dynamic graph databases, i.e. when graphs are added to or removed from the database. Third, we compared the indexes with regard to query workload changes. Fourth, we studied the scalability of our index vs. hybrid-indexes. Experimental results show that our index outperforms hybrid-indexes (i.e. indexes updated with one-pass) for query answering time in the case of social network databases, and is comparable

with these indexes for frequent and infrequent queries on chemical databases. Our graph-coarsening index can be updated up to 60 times faster in comparison to one-pass on dynamic graph databases. Moreover, our index is independent of the query workload for index update and is up to 15 times better after hybrid indexes are attuned to query workload for social network databases. This work is also published in 26th ACM International Conference on Information and Knowledge Management (CIKM) held in Singapore[18]."--Boise State University ScholarWorks.

Recently, the Semantic Web has gained huge popularity to address these challenges. Semantic web technologies have the opportunity to transform the way healthcare providers utilize technology to gain insights and knowledge from their data and make decisions. Both big data and semantic web technologies can complement each other to address the challenges and add intelligence to healthcare management systems. The aim of this book is to analyze the current status on how Semantic Web is used to solve the health data integration and interoperability problem, how it provides advanced data linking capabilities that can improve search and retrieval of medical data. There are chapters in the book which analyze the tools and approaches to semantic health data analysis and knowledge discovery. The book discusses the role of semantic technologies in extracting and transforming healthcare data before storing it in repositories. It also discusses different approaches for integrating heterogeneous healthcare data. To summarize, the book will help readers understand key concepts in semantic web applications for biomedical engineering and healthcare.

Graph databases provide a natural way of storing and querying graph data. In contrast to relational databases, queries over graph databases enable to refer directly to the graph structure of such graph data. For example, graph pattern matching can be employed to formulate queries over graph data. However, as for relational databases running complex queries can be very time-consuming and ruin the interactivity with the database. One possible approach to deal with this performance issue is to employ database views that consist of pre-computed answers to common and often stated queries. But to ensure that database views yield consistent query results in comparison with the data from which they are derived, these database views must be updated before queries make use of these database views. Such a maintenance of database views must be performed efficiently, otherwise the effort to create and maintain views may not pay off in comparison to processing the queries directly on the data from which the database views are derived. At the time of writing, graph databases do not support database views and are limited to graph indexes that index nodes and edges of the graph data for fast query evaluation, but do not enable to maintain pre-computed answers of complex queries over graph data. Moreover, the maintenance of database views in graph databases becomes even more challenging when negation and recursion have to be supported as in deductive relational databases. In this technical report, we present an approach for the efficient and scalable incremental graph view maintenance for deductive graph databases. The main concept of our approach is a generalized discrimination network that enables to model nested graph conditions including negative application conditions and recursion, which specify the content of graph views derived from graph data stored by graph databases. The discrimination network enables to automatically derive generic maintenance rules using graph transformations for maintaining graph views in case the graph data from which the graph views are derived change. We evaluate our approach in terms of a case study using multiple data sets derived from open source projects.

Computer and Information Sciences is a unique and comprehensive review of advanced technology and research in the field of Information Technology. It provides an up to date snapshot of research in Europe and the Far East (Hong Kong, Japan and China) in the most active areas of information technology, including Computer Vision, Data Engineering, Web Engineering, Internet Technologies, Bio-Informatics and System Performance Evaluation Methodologies.

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. This second edition includes new code samples and diagrams, using the latest Neo4j syntax, as well as information on new functionality. Learn how different organizations are using graph databases to outperform their competitors. With this book's data modeling, query, and code examples, you'll quickly be able to implement your own solution. Model data with the Cypher query language and property graph model Learn best practices and common pitfalls when modeling with graphs Plan and implement a graph database solution in test-driven fashion Explore real-world examples to learn how and why organizations use a graph database Understand common patterns and components of graph database architecture Use analytical techniques and algorithms to mine graph database information

The Easy, Common-Sense Guide to Solving Real Problems with NoSQL The Mere Mortals® tutorials have earned worldwide praise as the clearest, simplest way to master essential database technologies. Now, there's one for today's exciting new NoSQL databases. NoSQL for Mere Mortals guides you through solving real problems with NoSQL and achieving unprecedented scalability, cost efficiency, flexibility, and availability. Drawing on 20+ years of cutting-edge database experience, Dan Sullivan explains the advantages, use cases, and terminology associated with all four main categories of NoSQL databases: key-value, document, column family, and graph databases. For each, he introduces pragmatic best practices for building high-value applications. Through step-by-step examples, you'll discover how to choose the right database for each task, and use it the right way. Coverage includes --Getting started: What NoSQL databases are, how they differ from relational databases, when to use them, and when not to Data management principles and design criteria: Essential knowledge for creating any database solution, NoSQL or relational --Key-value databases: Gaining more utility from data structures --Document databases: Schemaless databases, normalization and denormalization, mutable documents, indexing, and design patterns --Column family databases: Google's BigTable design, table design, indexing, partitioning, and Big Data Graph databases: Graph/network modeling, design tips, query methods, and traps to avoid Whether you're a database developer, data modeler, database user, or student, learning NoSQL can open up immense new opportunities. As thousands of database professionals already know, For Mere Mortals is the fastest, easiest route to mastery.

Graph databases have become an indispensable tool for the analysis of linked data and interrelated data. As with any graph data representation, the need for using graph database systems emerge when they increase in size and complexity. Affiliated to those

needs, graph database benchmarks emerge to assess the performance of such systems in application scenarios, representative of use cases. I proposed graph database benchmarks based on the idea of social network application. The benchmark implements and proposes a data generator that synthetically generate graphs. Also a set of high level application queries on this application that model parts of the behavior of social network users. I also studied the graph database benchmarks results that affects the performance of graph database systems. I proposed graph database benchmarks in terms of data model, query workload, large datasets and usage scenarios. I discussed the characteristics of Neo4j and DEX graph databases to be included in the benchmarks. Also I studied, the characteristics of these graph databases queries that are important in the application of graph analysis.

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems. Learn how different organizations are using graph databases to outperform their competitors. With this book's data modeling, query, and code examples, you'll quickly be able to implement your own solution. Model data with the Cypher query language and property graph model Learn best practices and common pitfalls when modeling with graphs Plan and implement a graph database solution in test-driven fashion Explore real-world examples to learn how and why organizations use a graph database Understand common patterns and components of graph database architecture Use analytical techniques and algorithms to mine graph database information

Discover how graph databases can help you manage and query highly connected data. With this practical book, you'll learn how to design and implement a graph database that brings the power of graphs to bear on a broad range of problem domains. Whether you want to speed up your response to user queries or build a database that can adapt as your business evolves, this book shows you how to apply the schema-free graph model to real-world problems.

RDF-based knowledge graphs require additional formalisms to be fully context-aware, which is presented in this book. This book also provides a collection of provenance techniques and state-of-the-art metadata-enhanced, provenance-aware, knowledge graph-based representations across multiple application domains, in order to demonstrate how to combine graph-based data models and provenance representations. This is important to make statements authoritative, verifiable, and reproducible, such as in biomedical, pharmaceutical, and cybersecurity applications, where the data source and generator can be just as important as the data itself. Capturing provenance is critical to ensure sound experimental results and rigorously designed research studies for patient and drug safety, pathology reports, and medical evidence generation. Similarly, provenance is needed for cyberthreat intelligence dashboards and attack maps that aggregate and/or fuse heterogeneous data from disparate data sources to differentiate between unimportant online events and dangerous cyberattacks, which is demonstrated in this book. Without provenance, data reliability and trustworthiness might be limited, causing data reuse, trust, reproducibility and accountability issues. This book primarily targets researchers who utilize knowledge graphs in their methods and approaches (this includes researchers from a variety of domains, such as cybersecurity, eHealth, data science, Semantic Web, etc.). This book collects core facts for the state of the art in provenance approaches and techniques, complemented by a critical review of existing approaches. New research directions are also provided that combine data science and knowledge graphs, for an increasingly important research topic.

This book constitutes the refereed proceedings of the 15th European Workshop on Computer Performance Engineering, EPEW 2018, held in Paris, France, in October 2018. The 17 papers presented together with the abstracts of two invited talks in this volume were carefully reviewed and selected from 27 submissions. The papers presented at the workshop reflect the diversity of modern performance engineering, with topics ranging from advances in performanceengineering realm, including, dependability and security modeling, performance oriented model verification and testing, hardware and software systems case-studies, applications/extensions of queuing theory and network design

This book constitutes the revised selected papers of the 10th International Workshop on Information Search, Integration and Personalization, ISIP 2015, held in Grand Forks, ND, USA, in October 2015. The 8 revised full papers presented were carefully reviewed and selected from 26 submissions. The papers are organized in topical sections on modeling, querying and updating of information; information extraction; information visualization.

In scholarly digital editing, the established practice for semantically enriching digital texts is to add markup to a linear string of characters. Graph data-models provide an alternative approach, which is increasingly being given serious consideration. Labelled-property-graph databases, and the W3c's semantic web recommendation and associated standards (RDF and OWL) are powerful and flexible solutions to many of the problems that come with embedded markup. This volume explores the combination of scholarly digital editions, the graph data-model, and the semantic web from three perspectives: infrastructures and technologies, formal models, and projects and editions.

This book constitutes the proceedings of two events held at the CAiSE conference and relating to the areas of enterprise, business process and information systems modeling: The 19th International Conference on Business Process Modeling, Development and Support, BPMDS 2018, and the 23rd International Conference on Evaluation and Modeling Methods for Systems Analysis and Development, EMMSAD 2018. The conferences took place in Tallinn, Estonia, in June 2018. The 13 papers accepted for BPMDS were carefully reviewed and selected from 29 submissions; for EMMSAD 6 papers out of 13 submissions were accepted for publication. For BPMDS 2018, the papers were organized in topical sections as follows: context-awareness in business processes; automatic analysis of business processes; advanced approaches for business process modeling; evaluation of business process modeling techniques; an experience report on modeling collaborative processes. For EMMSAD 2018, the six related papers are listed without further sections.

Graphs have become one of the preferred ways to store structured data for various applications such as social network graphs, complex molecular structure, etc. Proliferation of graph databases has resulted in a growing need for effective querying methods to retrieve desired information. Querying has been widely studied in relational databases where the query optimizer finds a sequence of query execution steps (or plans) for efficient execution of the given query. Until now, most of the work on graph databases has concentrated on mining. For querying graph databases, users have to either learn a graph query language for posing their queries or use provided customized searches of specific substructures. Hence, there is a clear need for posing queries using graphs, consider alternative plans, and select a plan that can be processed efficiently on the graph database. In this thesis, we propose an approach to generate plans from a query using a cost-based approach that is tailored to the characteristics of the graph database. We collect metadata pertaining to the graph database and use cost estimates to evaluate the cost of execution of each plan. We use a branch and bound algorithm to limit the state space generated for identifying a good plan. Extensive experiments on different types of queries over two graph databases (IMDB and DBLP) are performed to validate our

approach. Subdue a graph mining algorithm has been modified to process a query plan instead of performing mining.

With coverage of the entire research process in social media, data collection and analysis on specific platforms, and innovative developments in the field, this handbook is the ultimate resource for those looking to tackle the challenges that come with doing research in this sphere.

This book constitutes the refereed proceedings of the 14th International Conference on Tests and Proofs, TAP 2020, held as part of the 4th World Congress on Formal Methods 2020, Bergen, Norway, in June 2020. The 7 regular papers, 1 short paper and 2 demonstration papers presented in this volume were carefully reviewed and selected from 209 submissions. The TAP conference promotes research in verification and formal methods that targets the interplay of proofs and testing: the advancement of techniques of each kind and their combination, with the ultimate goal of improving software and system dependability.

International Academic Conference on Management, Economics, Business and Marketing in Vienna, Austria 2016 (IAC-MEBM 2016), November 25 - 26, 2016

Why have developers at places like Facebook and Twitter increasingly turned to graph databases to manage their highly connected big data?

The short answer is that graphs offer superior speed and flexibility to get the job done. It's time you added skills in graph databases to your toolkit. In *Practical Neo4j*, database expert Greg Jordan guides you through the background and basics of graph databases and gets you quickly up and running with Neo4j, the most prominent graph database on the market today. Jordan walks you through the data modeling stages for projects such as social networks, recommendation engines, and geo-based applications. The book also dives into the configuration steps as well as the language options used to create your Neo4j-backed applications. Neo4j runs some of the largest connected datasets in the world, and developing with it offers you a fast, proven NoSQL database option. Besides those working for social media, database, and networking companies of all sizes, academics and researchers will find Neo4j a powerful research tool that can help connect large sets of diverse data and provide insights that would otherwise remain hidden. Using *Practical Neo4j*, you will learn how to harness that power and create elegant solutions that address complex data problems. This book: Explains the basics of graph databases Demonstrates how to configure and maintain Neo4j Shows how to import data into Neo4j from a variety of sources Provides a working example of a Neo4j-based application using an array of language of options including Java, .Net, PHP, Python, Spring, and Ruby As you'll discover, Neo4j offers a blend of simplicity and speed while allowing data relationships to maintain first-class status. That's one reason among many that such a wide range of industries and fields have turned to graph databases to analyze deep, dense relationships. After reading this book, you'll have a potent, elegant tool you can use to develop projects profitably and improve your career options.

This book is simply the introduction to data modeling using a simple, straightforward scenario. There are plenty of opportunities throughout the upcoming guides to practice modeling domains and analyzing changes to the model that might need to be made. Neo4j is an open source NoSQL graph database. It is a fully transactional database (ACID) that stores data structured as graphs consisting of nodes, connected by relationships. Inspired by the structure of the real world, it allows for high query performance on complex data, while remaining intuitive and simple for the developer. Using this book, you'll get to learn the theory of graph database and how to use Neo4j to build up recommendations, relationships, and calculate the shortest route between two locations. With example data models, best practices, use-cases, and an application putting everything together, this book will give you everything you need to really get started with Neo4j. Starting with a brief introduction to graph theory, this book will show you the advantages of using graph databases along with data modeling techniques for graph databases. You'll gain practical hands-on experience with commonly used and lesser known features for updating graph store with Neo4j's Cypher query language. This book includes a lot of background information, helps you grasp the fundamental concepts behind this radical new way of dealing with connected data, and will give you lots of examples of use cases and environments where a graph database would be a great interest.

Global growth is in low gear, and the drivers of activity are changing. These dynamics raise new policy challenges.

Advanced economies are growing again but must continue financial sector repair, pursue fiscal consolidation, and spur job growth. Emerging market economies face the dual challenges of slowing growth and tighter global financial conditions. This issue of the *World Economic Outlook* examines the potential spillovers from these transitions and the appropriate policy responses. Chapter 3 explores how output comovements are influenced by policy and financial shocks, growth surprises, and other linkages. Chapter 4 assesses why certain emerging market economies were able to avoid the classical boom-and-bust cycle in the face of volatile capital flows during the global financial crisis.

We are now in an era where the technology has rapidly become democratized. The restrictions of relational databases to address the requirements of contemporary application domains, such as semantic web and social networking, where data has an inherited graph structure underlying in it, leads to the development of new technology called Graph Databases. Graph databases can be defined as those in which data structures for the schema and instances are modeled as graphs and data manipulation is expressed by graph-traversal operations. On the other hand, the Complexity involved in the relational model to process large queries quickly involving complex join operations leads to the development of the Join-Core. The Join Core consists of a set of tables that store the relationships of data. With join core, no relations or intermediate results need to be retrieved, generated, or transferred, only query results need to be transferred over the networks. In this study, an overview and comparison of current graph database models like AllegroGraph, FlockDB, Neo4j, Sones, DEX etc. is presented. The comparison shows that Neo4j is the most popular and highly recommended among all graph databases. Also the experimental results of Join Core against Neo4j graph database shows that join core performance is more efficient than Neo4j when the query has complex join operations involved in it.

"We start off with an overview to graph modeling, followed by an explanation of graph databases and their roles within our organizations. Learn the top use cases for graph databases, along with best practices in graph modeling. Neo4j is used as an example."--Resource description page.

Neo4j is a graph database that allows us to model our data as a graph and find solutions to complex real-world problems that are difficult to solve using any other type of database. This book is designed and presented to help you understand the intricacies of modeling a graph for any practical domain. The book starts with an example of a graph problem and then introduces you to modeling non-graph problems using Neo4j. Throughout the book, you will discover design choices and trade-offs, and understand how and when to use Neo4j. Starting with a brief introduction to graph theory, this book will show you the advantages of using graph databases along with data modeling techniques for graph databases. You'll gain practical hands-on experience with commonly used and lesser known features for updating graph store with Neo4j's Cypher query language. This book includes a lot of background information, helps you grasp the fundamental concepts

practitioners, researchers and academics, as well as undergraduate and postgraduate students.

The field of health is an increasingly complex and technical one; and an area in which a more multidisciplinary approach would undoubtedly be beneficial in many ways. This book presents papers from the conference 'Health – Exploring Complexity: An Interdisciplinary Systems Approach', held in Munich, Germany, from August 28th to September 2nd 2016. This joint conference unites the conferences of the German Association for Medical Informatics, Biometry and Epidemiology (GMDS), the German Society for Epidemiology (DGEpi), the International Epidemiological Association - European Region, and the European Federation for Medical Informatics (EFMI). These societies already have long-standing experience of integrating the disciplines of medical informatics, biometry, epidemiology and health data management. The book contains over 160 papers, and is divided into 14 sections covering subject areas such as: health and clinical information systems; eHealth and telemedicine; big data and advanced analytics; and evidence-based health informatics, evaluation and education, among many others. The book will be of value to all those working in the field of health and interested in finding new ways to enable the collaboration of different scientific disciplines and the establishment of comprehensive methodological approaches.

Graph Databases in Action teaches readers everything they need to know to begin building and running applications powered by graph databases. Right off the bat, seasoned graph database experts introduce readers to just enough graph theory, the graph database ecosystem, and a variety of datastores. They also explore modelling basics in action with real-world examples, then go hands-on with querying, coding traversals, parsing results, and other essential tasks as readers build their own graph-backed social network app complete with a recommendation engine! Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications.

[Copyright: 5178e1fae993ffdc824763ab81cc290c](https://www.manning.com/books/graph-databases-in-action)