

Top Three Data Modeling Trends

Steve Hoberman

me@stevehoberman.com
TechnicsPub.com
DataModelingZone.com

data modeling zone

10TH EDITION
DATA
MODELING MASTER CLASS
TRAINING MANUAL
Steve Hoberman's Best Practice Approach to
Designing and Applying Data Models
Using Advanced Modeling Software
STEVE HOBERMAN

1

Data Modeling Trends

- Trends
 - Mainstream NoSQL projects (“Modernizing”)
 - Knowledge graphs getting the spotlight
 - AI growing, growing, growing
- More connecting and less explanatory
- Sharing over lecturing
- Your insights as well please!
- Learn at least three new things!



2

Data Modeling Trend #1

- Mainstream NoSQL projects (“Modernizing”)
- Impact
 - Traditional modelers becoming NoSQL modelers/enablers (or not knowing about these projects)
 - “Embrace NoSQL”
- Response
 - Retaining, Redefining, Recommending (Camera settings, CDM/LDM, patterns)
 - Learning (DMZ)



3

Data Model Settings

A data model contains settings, similar to a camera...



Masterpiece

Understand the need for the picture and then choose the most appropriate setting values

4

Data Model Settings

Mode

Timer

Scope

...and compose!

5

Scope

(What is the scope of the initiative?)

Scope

Bounded Context / Domain


Eric Evans

Project	<ul style="list-style-type: none"> Fraud detection for auto insurance Northeast Products Analytics London hotel survey application
Program	<ul style="list-style-type: none"> Fraud detection for all LOB Enterprise Products Analytics Hotel chain survey application

7

Focus

(Which is more important, flexibility or simplicity?)



Focus

Only setting which can be at a subject level within a model

Abstract (flexible)	<ul style="list-style-type: none"> Party Event Intellectual Unit
Concrete (simple)	<ul style="list-style-type: none"> Customer Order Book

9

Timer

(What do you do? What do you want to do?)



Timer




Now	Legal entities can cross countries.
Later	Legal entities cannot cross countries.

11

Filter

(Are we forward or reverse engineering?)




Business (forward)	<ul style="list-style-type: none"> The business calls it a 'classification'. What patterns are of interest? I need a 360 degree view of customer.
Application (reverse)	<ul style="list-style-type: none"> SAP calls it a 'characteristic'. What connections exist in this data set? I need to extract fields from this JSON file.

13

Mode

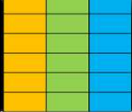

(What is the purpose of the application?)



RDBMS	<ul style="list-style-type: none"> Benefit: Precisely representing data through sets Captures: Business rules <i>constraining</i> a business process Primary use case: Operational (OLTP) Modeling: Relational <p>A Customer must own at least one Account.</p>
	<ul style="list-style-type: none"> Benefit: Precisely representing how data will be analyzed Captures: Business questions <i>analyzing</i> a business process Primary use case: Analytics (OLAP) Modeling: Dimensional <p>How much revenue did we generate in fees by Date, Region, and Product? Also want to see by Month and Year...</p>
NoSQL	<ul style="list-style-type: none"> Benefit: Precisely representing how data will be received and accessed Captures: Access paths <i>providing insights</i> into a business process Primary use case: Discovery Modeling: Query <p>Which customers own a checking account that generated over \$10,000 in fees this year, own at least one cat, and live within 500 miles of New York City?</p>


15

Mode (Purpose?): RDBMS


```

CREATE TABLE Book (
  bookTitleName VARCHAR2 (50) NOT NULL,
  bookDescription VARCHAR2 (255) NOT NULL,
  bookPageCount NUMBER (1, 0) NOT NULL,
  bookRetailPriceAmount NUMBER (1, 0) NOT NULL,
  bookAvailableIndicator CHAR (1) NOT NULL,
  bookPublicationDate DATE NOT NULL,
  bookCover BINARY_DOUBLE NOT NULL,
  CONSTRAINT PK1 PRIMARY KEY (bookTitleName))
  
```

bookTitleName	bookDescription	bookPublicationDate	bookRetailPriceAmount	bookPageCount	bookAvailableIndicator	bookCover
Building the Data Lakehouse	The data lakehouse is the next generation of the data warehouse and data lake, designed to meet today's complex and ever-changing analytics, machine learning, and data science requirements.	October 1, 2021	\$49.95	254	True	

16

Mode (Purpose?): NoSQL



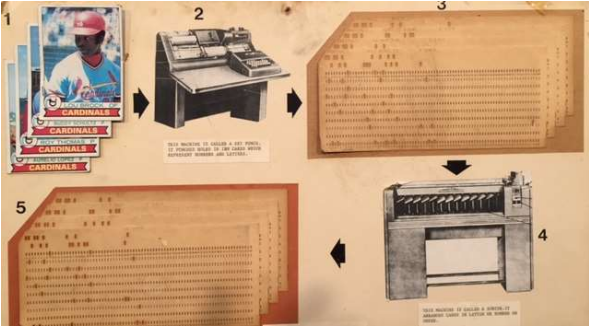
- NoSQL = Any database that is not a RDBMS
- Nothing to do with query languages!
- Performance-driven and reactionary
- First NoSQL assignment in 1979
- “Form follows function”
- A better name?

NoRDBMS?

NoSets?

NoJoins?

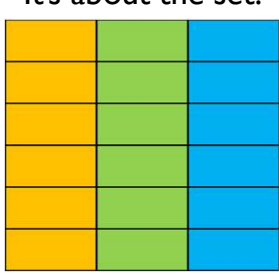
NoNew?



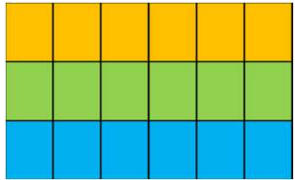
17

Each type of database sees data differently

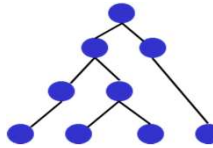
RDBMS
Set Theory (1870)
Codd (1970)
System R (1974)
"It's about the set."



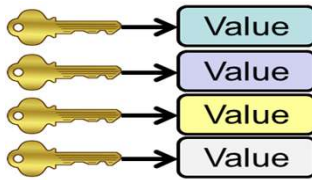
Column
TAXIR (1969), Redshift,
Hbase, Cassandra
"It's about the attribute."



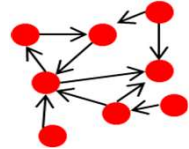
Document
Hierarchy Theory (1880)
Lotus Notes (1989)
MongoDB, MarkLogic,
RethinkDB, Couchbase
"It's about the doc."



Key-Value
DynamoDB (2007), Redis, Oracle
NoSQL, Couchbase
"It's about the value."

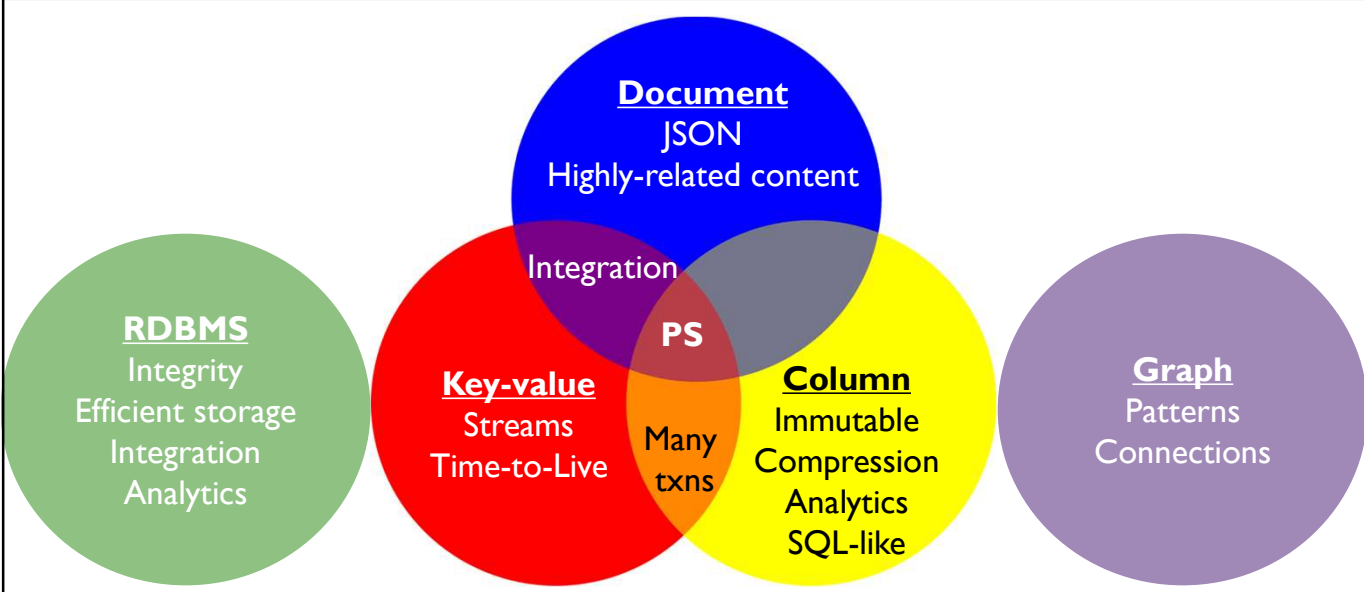


Graph
Graph Theory (1735)
Neo4j, AllegroGraph
"It's about the pattern."



18



When to use each type of database

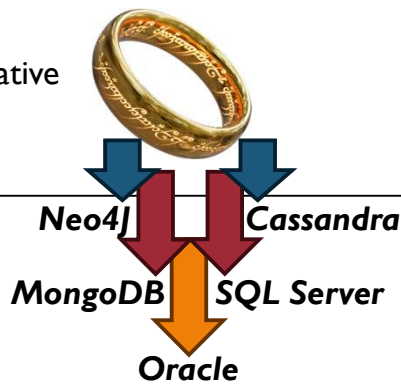


PS = Performance and Scalability

19

Align > Refine > Design Overview

<p><u>Align (Business Terms - BTM)</u> Common business vocabulary Important terms, their relationships, and definitions “This is what we need to photograph.”</p>	<p>“Which are your pets?” “Do you like landscapes or portraits?” “Action or posed shots?”</p>
<p><u>Refine (Logical - LDM)</u> What is needed for a specific initiative Independent of technology “We’ve taken the picture.”</p>	
<p><u>Design (Physical - PDM)</u> Instantiation (schema) Compromised for technology “We can print it on a mug.”</p>	



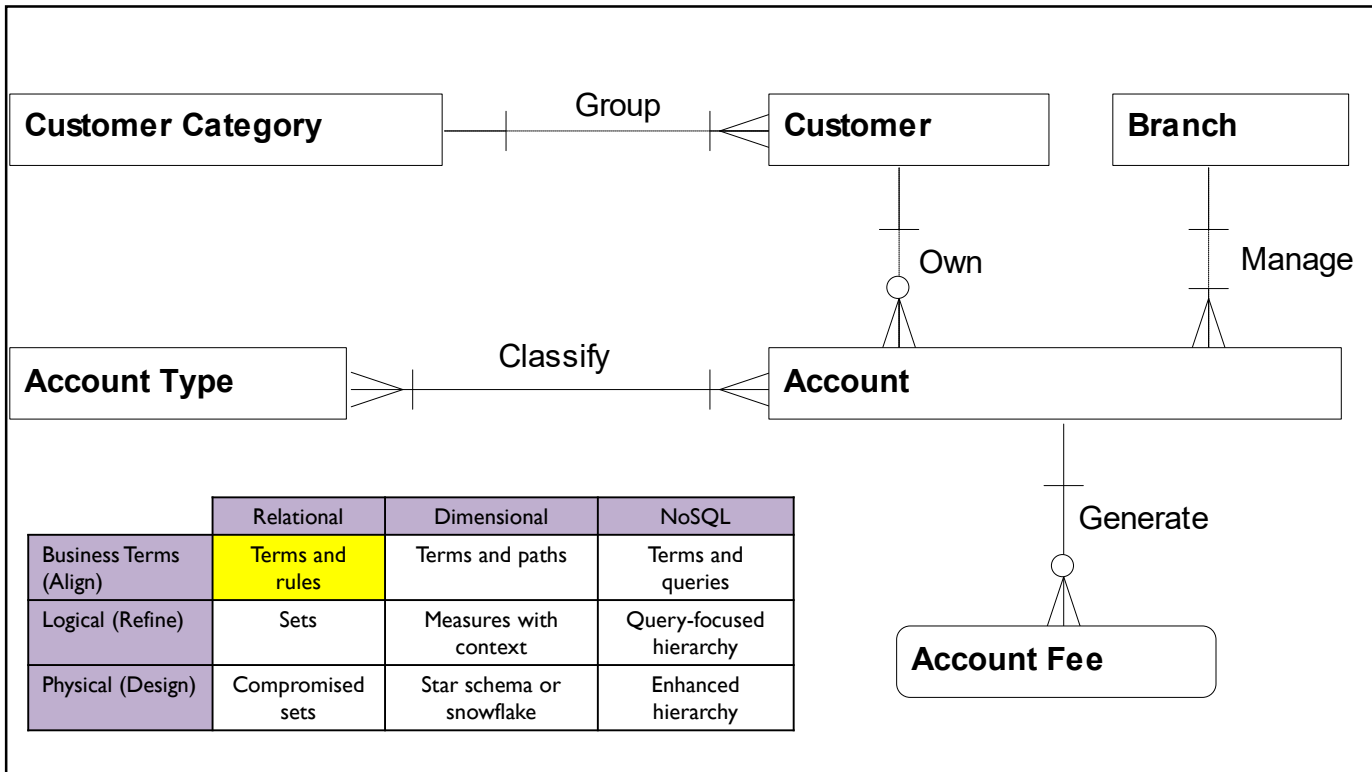
20

★ Overview

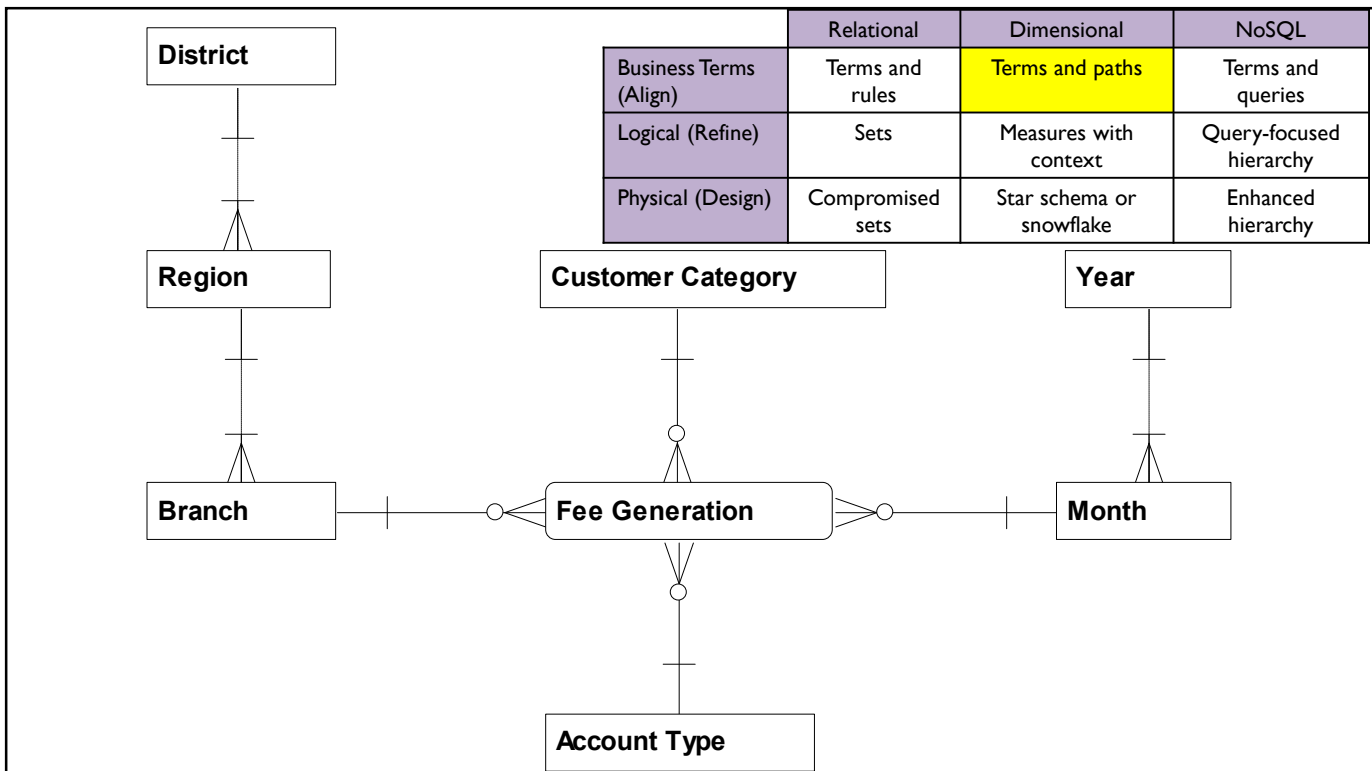
	Relational	Dimensional	NoSQL
Business Terms (Align)	Terms and rules	Terms and paths	Terms and queries
Logical (Refine)	Sets	Measures with context	Query-focused hierarchy*
Physical (Design)	Compromised sets	Star schema or snowflake	Enhanced hierarchy*

*Exceptions including graph

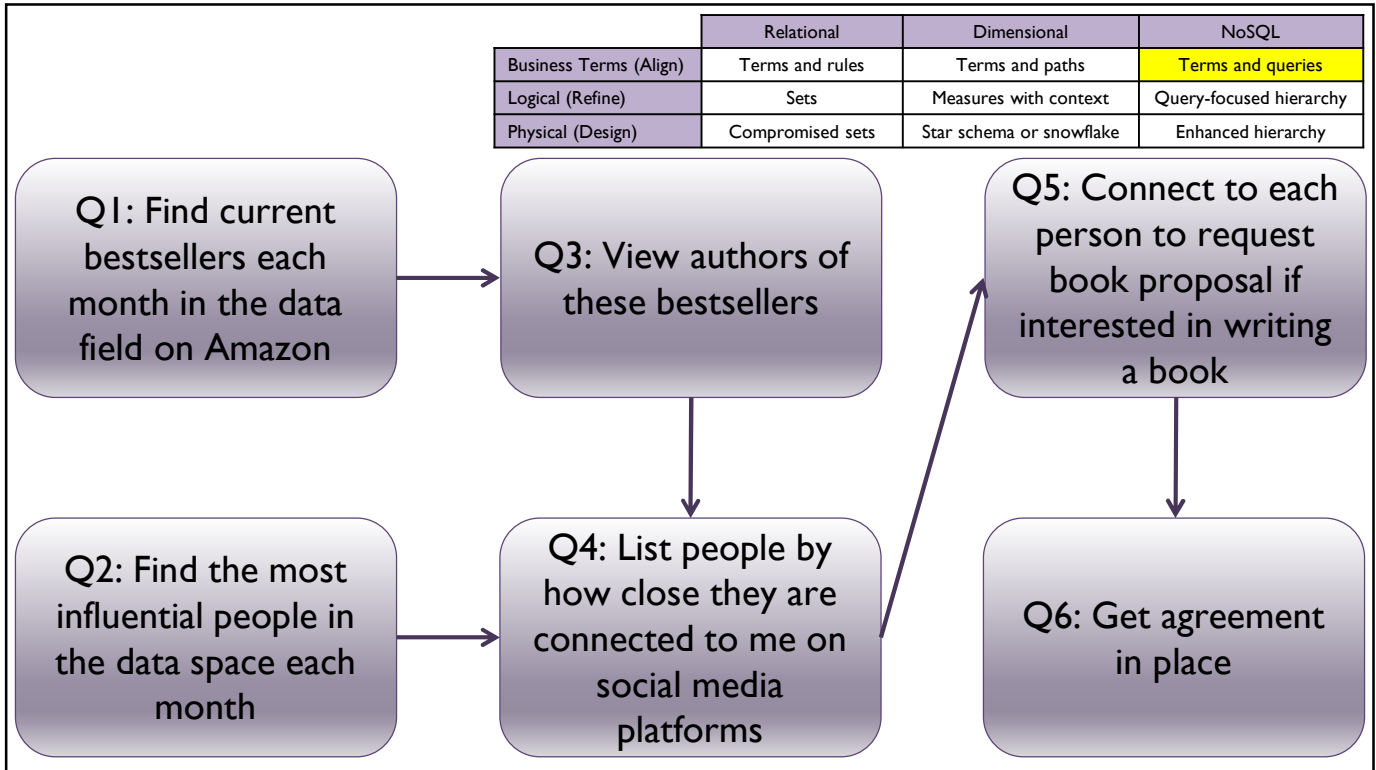
21



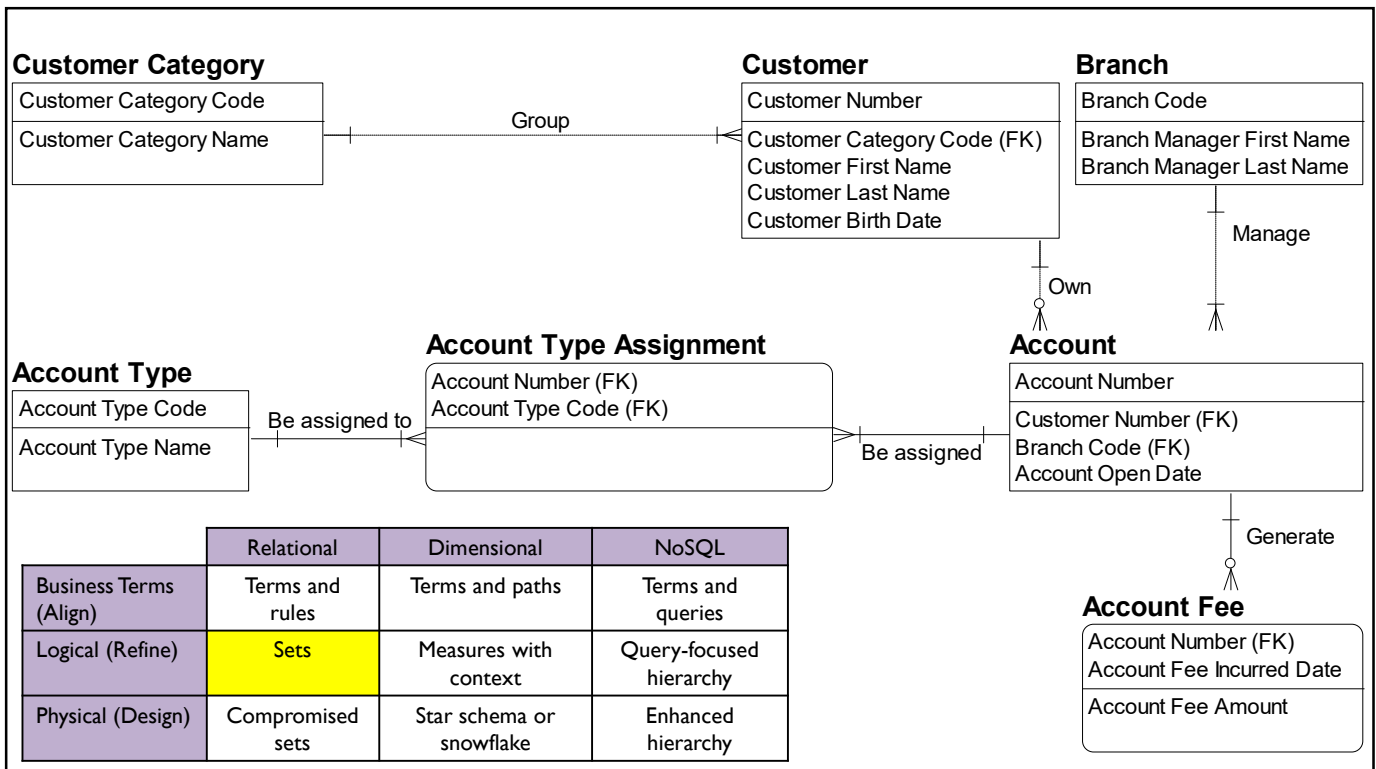
22



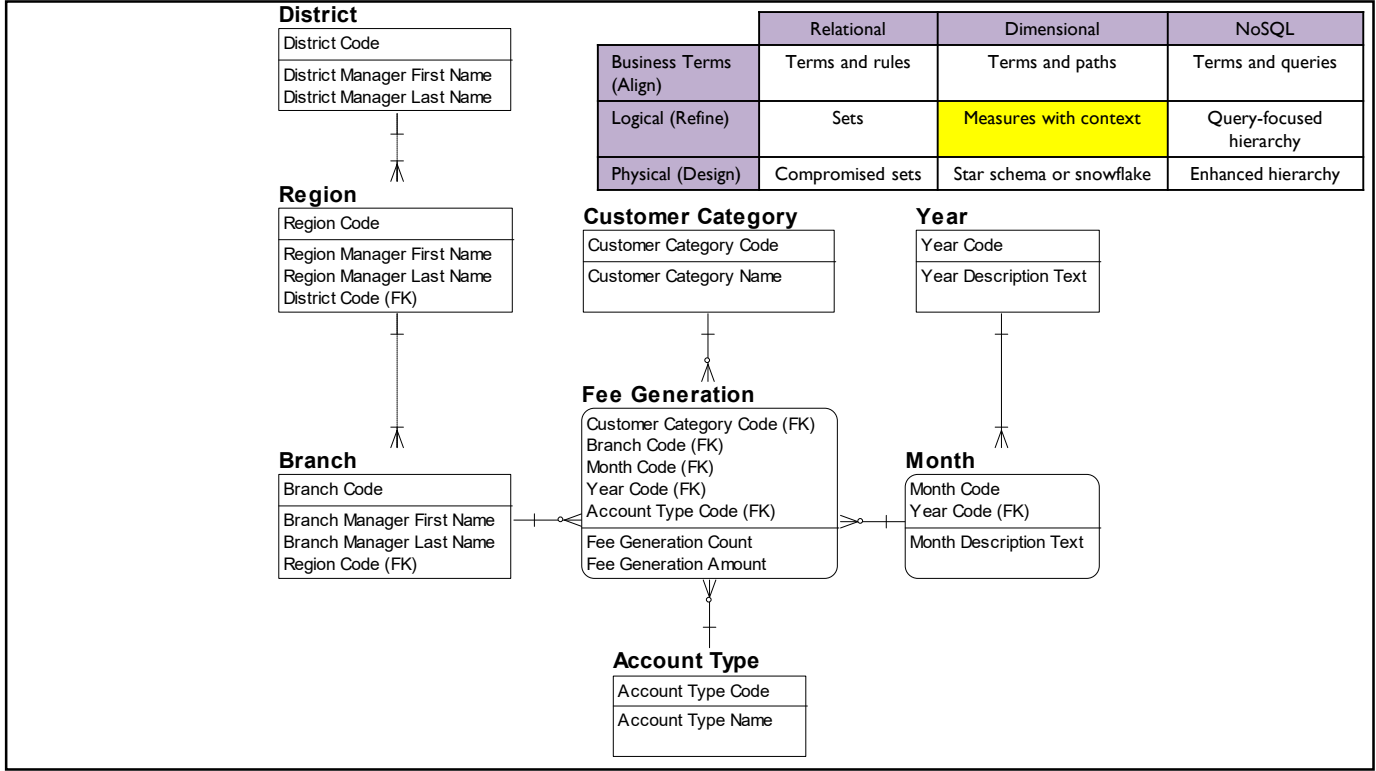
23



24



25



26

Most Potential to Produce a Bestseller (Q1, Q2, Q3, Q4)			
Profile ID	pk	int32	*
Person Name		vchar	*
LinkedIn URL		vchar	
LinkedIn Connection Count		num	
Person Accomplishment		arr	
[0] Person Accomplishment Description		vchar	
Person Skill		arr	
[0] Person Skill Name		vchar	
Person Experience		arr	
[0] Person Experience Description		vchar	
Person Country		arr	
[0] Person Country Name		vchar	
Person Connection		arr	
[0] Connection Name		vchar	
Author of Books		arr	
[0] Title Name		vchar	
[1] Subtitle Name		vchar	
[2] Title Price Amount		num	
[3] Amazon URL		vchar	
[4] Amazon Thumbnail		vchar	
Category		arr	
[0] Category Name		vchar	

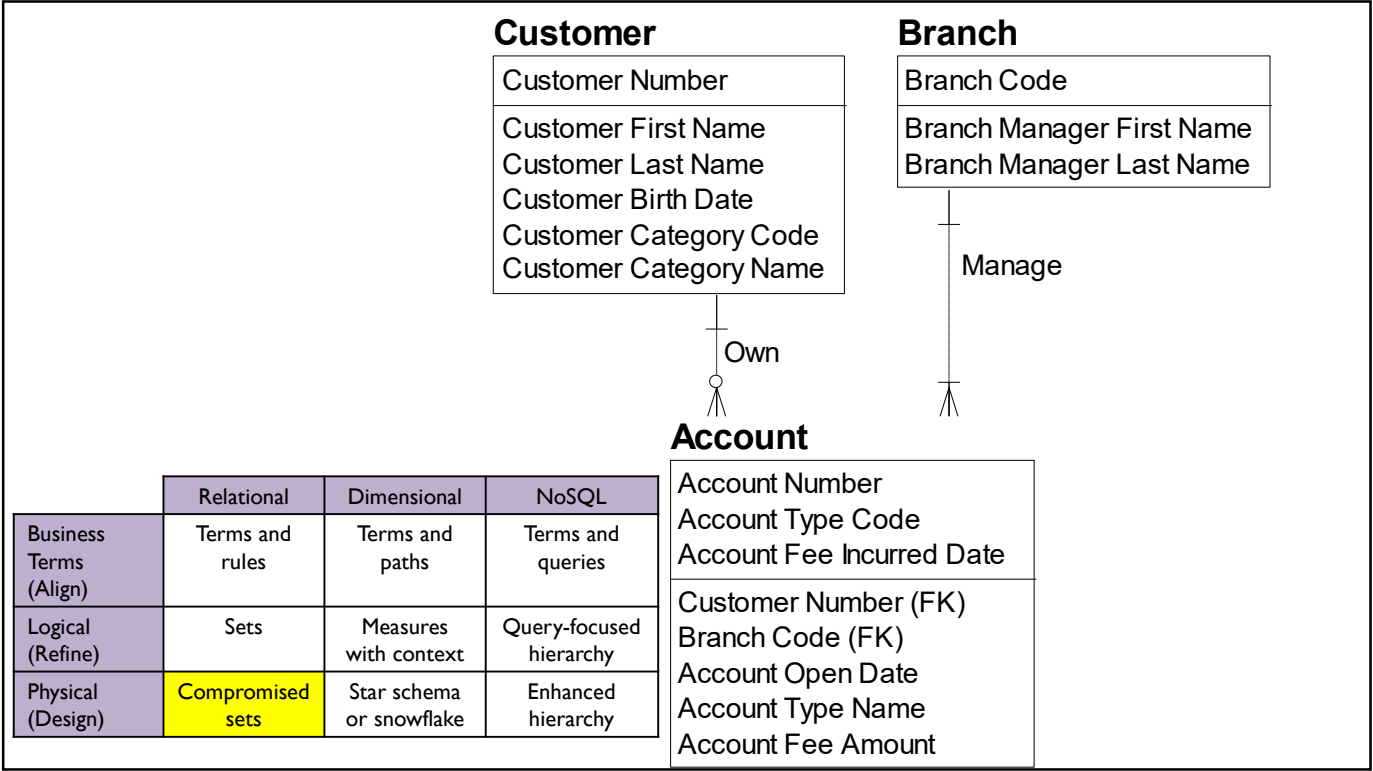
Book Proposal (Q5)			
Author Name	pk	vchar	*
Book Title	pk	vchar	*
Book Subtitle Name		vchar	*
Author Phone Number		vchar	
Author Email Address		vchar	
Author Address		vchar	
Estimated Page Count		num	
Estimated Manuscript Ready for Editing Date		date	
Back Cover Text		vchar	
Topic		arr	
[0] Topic Name		vchar	
Reason for Writing Book		vchar	
Competition		arr	
[0] Book Title Name		vchar	

Book Agreement (Q6)			
Author Name	pk	vchar	*
Book Title	pk	vchar	*
Edition Number	pk	vchar	*
Book Subtitle Name		vchar	*
Author Address		vchar	
Country Name		vchar	
Manuscript Due Date		date	
Book Royalty Percent		num	
Author Signature		bin	
Author Signature Date		date	
Fully Signed Author Document		bin	

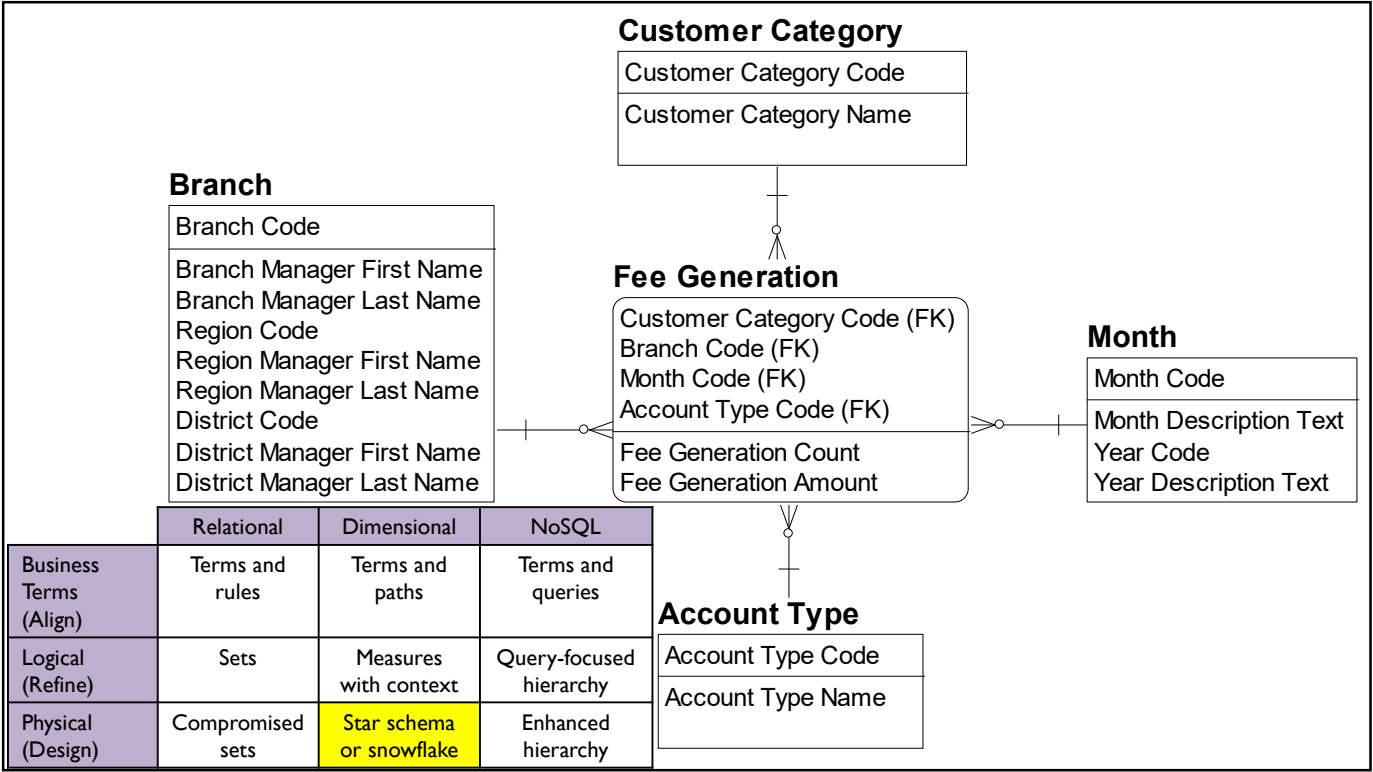
Polyglot model

	Relational	Dimensional	NoSQL
Business Terms (Align)	Terms and rules	Terms and paths	Terms and queries
Logical (Refine)	Sets	Measures with context	Query-focused hierarchy
Physical (Design)	Compromised sets	Star schema or snowflake	Enhanced hierarchy

27



28



29

DynamoDB

Business Terms (Align)	Relational	Dimensional	NoSQL
Terms and rules	Terms and rules	Terms and paths	Terms and queries
Logical (Refine)	Sets	Measures with context	Query-focused hierarchy
Physical (Design)	Compromised sets	Star schema or snowflake	Enhanced hierarchy

Most Potential to Produce a Bestseller (Q1, Q2, Q3, Q4)			
Profile ID	pk	N	*
Person Name		S	*
LinkedIn URL		S	
LinkedIn Connection Count		N	(12)
Person Accomplishment		L	
[0] Person Accomplishment Description		S	
Person Skill		L	
[0] Person Skill Name		S	
Person Experience		L	
[0] Person Experience Description		S	
Person Country		L	
[0] Person Country Name		S	
Person Connection		L	
[0] Connection Name		S	
Author of Books		L	(11)
[0] Title Name		S	
[1] Subtitle Name		S	
[2] Title Price Amount		N	
[3] Amazon URL		S	
[4] Amazon Thumbnail		S	
[5] Category		L	
[0] Category Name		S	

Book Proposal (Q5)			
Author Name	pk	S	*
Book Title	pk	S	*
Book Subtitle Name		S	*
Author Phone Number		S	
Author Email Address		S	
Author Address		S	
Estimated Page Count		N	
Estimated Manuscript Ready for Editing Date		S	
Back Cover Text		S	
Topic		L	
[0] Topic Name		S	
Reason for Writing Book		S	
Competition		L	
[0] Book Title Name		S	

Book Agreement (Q6)			
Author Name	pk	S	*
Book Title	pk	S	*
Edition Number	pk	S	*
Book Subtitle Name		S	*
Author Address		S	
Country Name		S	
Manuscript Due Date		S	
Book Royalty Percent		N	
Author Signature		B	
Author Signature Date		S	
Fully Signed Author Document		B	

Updated with the "secret sauce"

30

Cassandra

Business Terms (Align)	Relational	Dimensional	NoSQL
Terms and rules	Terms and rules	Terms and paths	Terms and queries
Logical (Refine)	Sets	Measures with context	Query-focused hierarchy
Physical (Design)	Compromised sets	Star schema or snowflake	Enhanced hierarchy

Most Potential to Produce a Bestseller (Q1, Q2, Q3, Q4)			
Profile ID	pk	integer	*
Person Name		varchar	*
LinkedIn URL		varchar	
LinkedIn Connection Count		integer	(2 2.1)
Person Accomplishment		list<str>	
[0] Person Accomplishment Description		varchar	
Person Skill		list<str>	
[0] Person Skill Name		varchar	
Person Experience		list<str>	
[0] Person Experience Description		varchar	
Person Country		list<str>	
[0] Person Country Name		varchar	
Person Connection		list<str>	
[0] Connection Name		varchar	
Author of Books		list<str>	(2 1.1)
[0] Title Name		varchar	
[1] Subtitle Name		varchar	
[2] Title Price Amount		integer	
[3] Amazon URL		varchar	
[4] Amazon Thumbnail		varchar	
[5] Category		list<str>	
[0] Category Name		varchar	

Book Proposal (Q5)			
Author Name	pk	varchar	*
Book Title	pk	varchar	*
Book Subtitle Name		varchar	*
Author Phone Number		varchar	
Author Email Address		varchar	
Author Address		varchar	
Estimated Page Count		integer	
Estimated Manuscript Ready for Editing Date		dt	
Back Cover Text		varchar	
Topic		list<str>	
[0] Topic Name		varchar	
Reason for Writing Book		varchar	
Competition		list<str>	
[0] Book Title Name		varchar	

Book Agreement (Q6)			
Author Name	pk	varchar	*
Book Title	pk	varchar	*
Edition Number	pk	varchar	*
Book Subtitle Name		varchar	*
Author Address		varchar	
Country Name		varchar	
Manuscript Due Date		dt	
Book Royalty Percent		integer	
Author Signature		blb	
Author Signature Date		dt	
Fully Signed Author Document		blb	

Updated with the "secret sauce"

31

MongoDB

Most Potential to Produce a Bestseller (Q1, Q2, Q3, Q4)

Profile ID	pk	int32	*
Person Name		str	*
LinkedIn URL		str	
LinkedIn Connection Count		num	(12.1)
Person Accomplishment		arr	
[0] Person Accomplishment Description		str	
Person Skill		arr	
[0] Person Skill Name		str	
Person Experience		arr	
[0] Person Experience Description		str	
Person Country		arr	
[0] Person Country Name		str	
Person Connection		arr	
[0] Connection Name		str	
Author of Books		arr	(11.1)
[0] Title Name		str	
[1] Subtitle Name		str	
[2] Title Price Amount		num	
[3] Amazon URL		str	
[4] Amazon Thumbnail		str	
[5] Category		arr	
[0] Category Name		str	

Book Proposal (Q5)

Author Name	pk	str	*
Book Title	pk	str	*
Book Subtitle Name		str	*
Author Phone Number		str	
Author Email Address		str	
Author Address		str	
Estimated Page Count		num	
Estimated Manuscript Ready for Editing Date		date	
Back Cover Text		str	
Topic		arr	
[0] Topic Name		str	
Reason for Writing Book		str	
Competition		arr	
[0] Book Title Name		str	

	Relational	Dimensional	NoSQL
Business Terms (Align)	Terms and rules	Terms and paths	Terms and queries
Logical (Refine)	Sets	Measures with context	Query-focused hierarchy
Physical (Design)	Compromised sets	Star schema or snowflake	Enhanced hierarchy

Updated with the "secret sauce"

32


Common data modeling patterns

All the benefits of abstraction

Designs created faster

Development time reduced

Improved consistency across models




THE NIMBLE ELEPHANT


AGILE DELIVERY OF DATA MODELS USING A PATTERN-BASED APPROACH

JOHN GILES
Foreword by Scott Ambler


→



Elementary

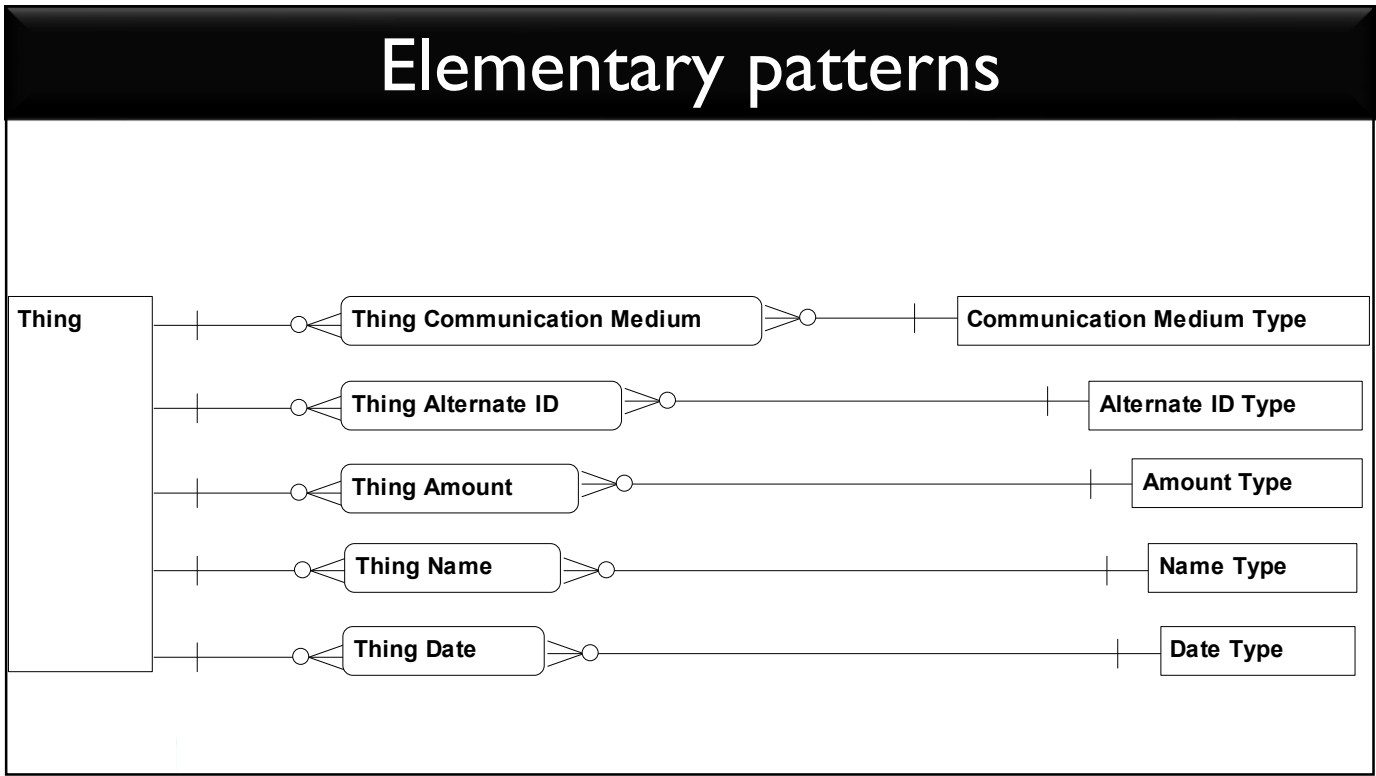


Assembly

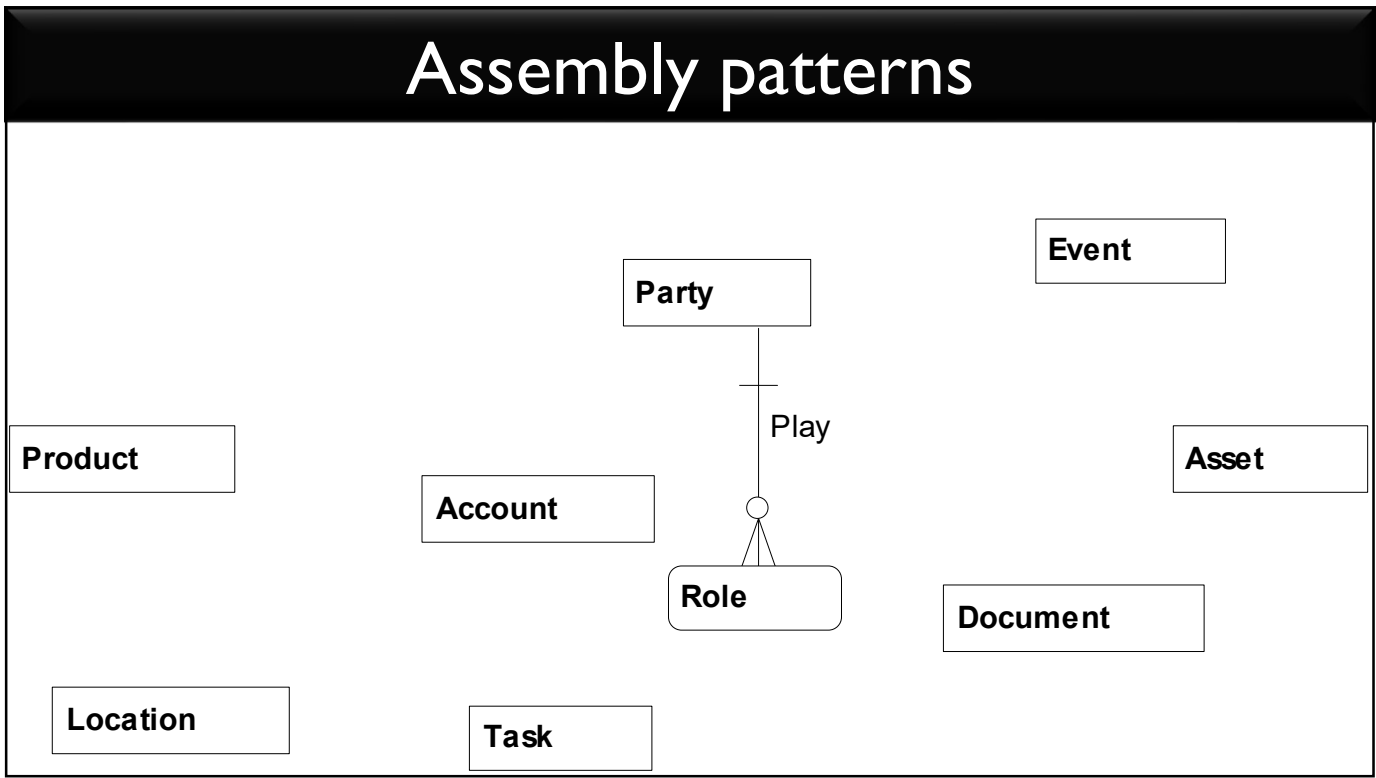


Integration

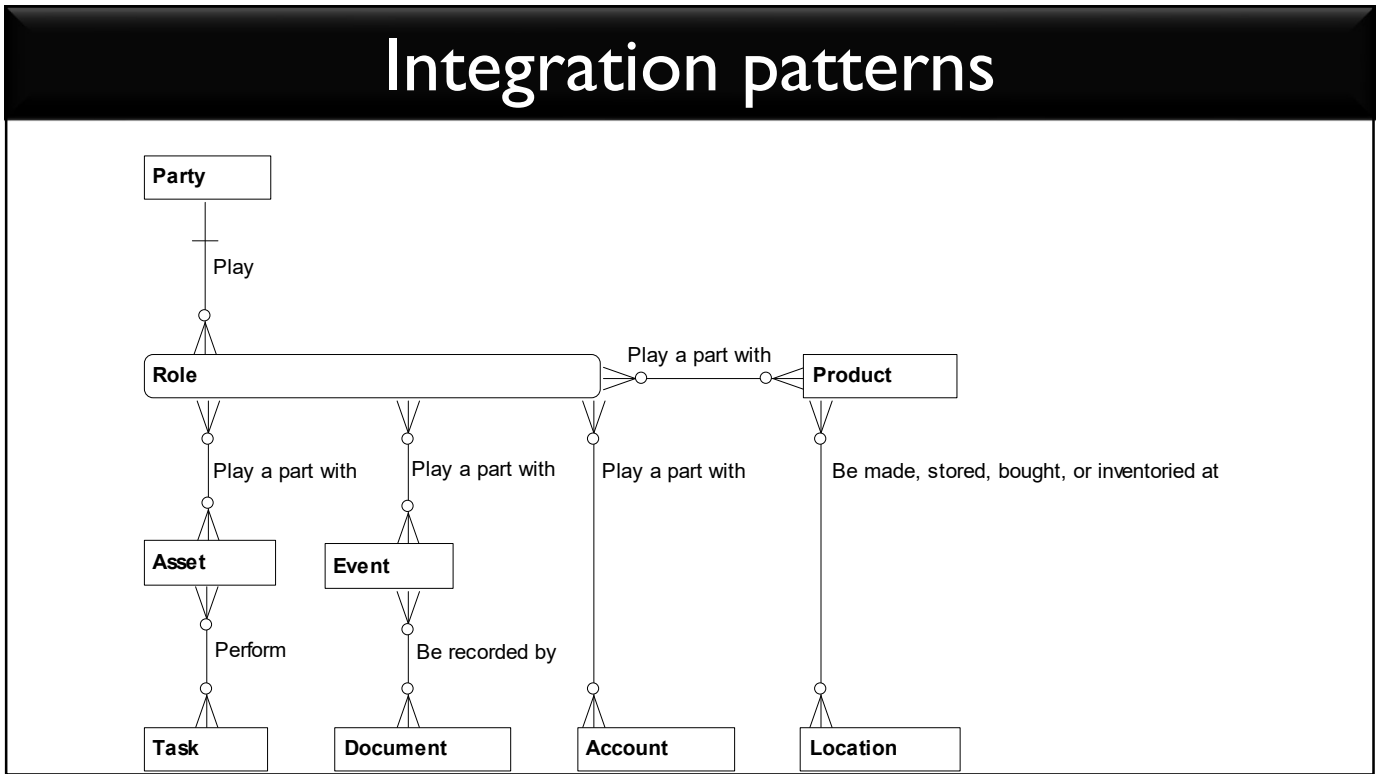
33



34



35



36

Common NoSQL patterns

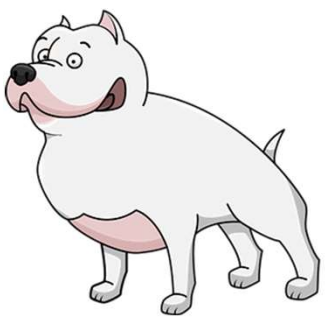
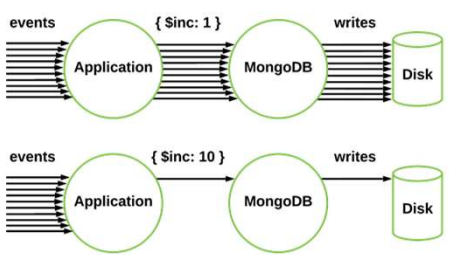
Category	Patterns
Computation	<ul style="list-style-type: none"> • Approximation
Grouping	<ul style="list-style-type: none"> • Bucket
Lifecycle	<ul style="list-style-type: none"> • Document Versioning • Envelope
Polymorphism	<ul style="list-style-type: none"> • Inheritance

And many more!!

37

The Approximation Pattern

- What is the population of the Earth?
- How many people have visited our website today?
- When the cost of perfect numbers is not worth the effort
- Too many writes and exact numbers are not needed or not practical



Approximation			
_id	pk	str	*
pet_name		str	*
web_page_views		str	

```

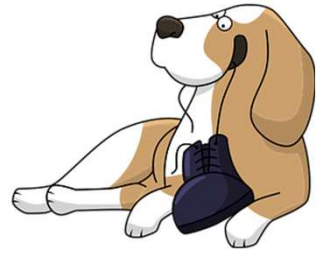
if random(0..99) == 0:
    increment = 100
else:
    increment = 0
  
```

38

The Bucket Pattern

- Repeating Attributes technique + optionally summarization
- Fixed arrays to avoid references
- One-to-many that can't be embedded
- IoT, DW

Bucket			
_id	pk	doc	*
pet_id		str	
month		date	
interactions		arr	
[0]		doc	
ts		date	
userid		int32	

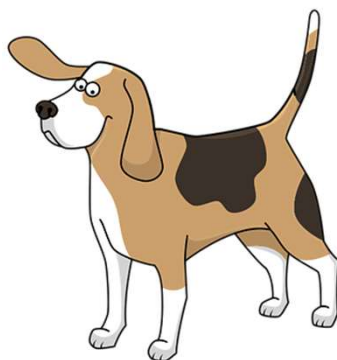


Account
Account Number
Customer Number (FK)
Branch Code (FK)
Account Open Date
Account Type Code 1
Account Type Name 1
Account Type Code 2
Account Type Name 2
Account Type Code 3
Account Type Name 3

39

The Document Versioning Pattern

- SCDs for documents
- Keep history
- Financial, insurance, and legal apps



DocVersioningLatest			
customer_id	pk	str	*
revision		int32	*
ts		date	*
name		str	
insured_items		arr	
[0] home		doc	
type		str	*
address		str	
[1] life		doc	
type		str	*

40

The Envelope Pattern

- Separate data intended for consumption from data intended to optimize the power and flexibility of the database and application.
- Dat lineage, integration

Envelope			
_id	pk	str	*
header		doc	
schema_version		num	
docRevision		num	
creation_ts		date	
last_update_ts		date	
created_by		str	
provenance		doc	
source		str	
lineage		str	
harmonization		doc	
zipcode		str	
phone		str	
related_to		arr	
[0]		str	
payload		doc	
...		str	

Envelope wraps metadata like version numbers, timestamps, lineage, or other data used for indexing or other database operations.

Payload contains the actual data intended to be consumed by the application.

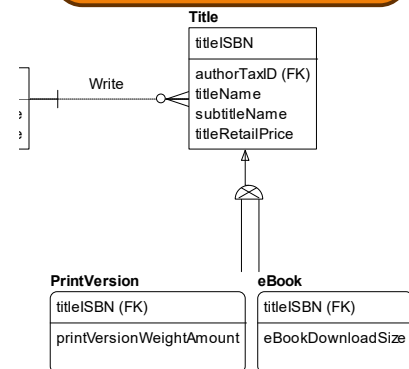
41

The Inheritance Pattern

- Put common fields in the document's root and isolate specificities in a subdocument

inheritance-Product			
_id	pk	str	*
object_type		str	*
price		dec128	
oneOf		ch	
[0]		sub	
book		doc	
author		str	
publisher		str	
[1]		sub	
shirt		doc	
fabric		str	
size		doc	
width		str	
height		str	
depth		str	
code		str	

Recall "Rolling Up",
but document style



42

Education: DMZ Sessions

- A Practical Use Case for NoSQL Data Modeling at American Express. Eve Danoff, American Express
- MongoDB Schema Design Patterns. Daniel Coupal, MongoDB
- Use JSON models to give your logical models life as a semantic layer. Rob Garrison, Compassion International
- Oracle 23c Duality Views: expose data stored in relational tables as JSON documents and get the best of both worlds! Pascal Desmarests, Hackolade and Beda Hammerschmidt, Oracle
- LDM + NoSQL? Mini-Hackathon. Pascal Desmarests and Marco Wobben



43

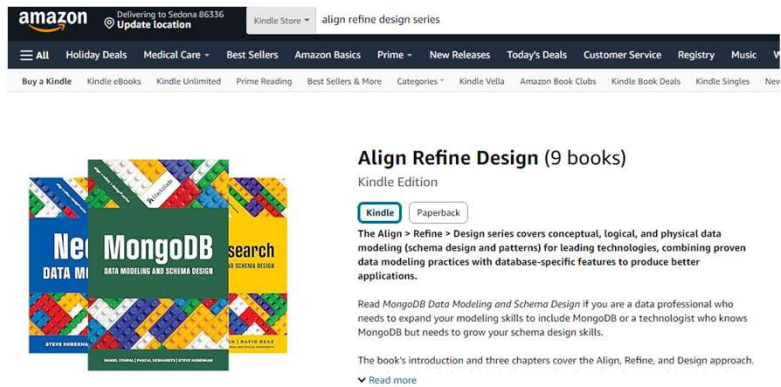
Publishing

- Align > Refine > Design Series
- <https://technicspub.com/data-modeling-design/>

78 • MongoDB Data Modeling and Schema Design

of normalization is to correctly organize attributes into sets. Also, note that the normalized model is built according to the properties of the data and not built according to how the data is being used.

Dimensional models are built to answer specific business questions with ease, and NoSQL models are built to answer queries and identify patterns with ease. The relational model is the only model focused on the intrinsic properties of the data and not usage.



44

Data Modeling Trend #1

- Mainstream NoSQL projects (“Modernizing”)
- Impact
 - Traditional modelers becoming NoSQL modelers/enablers (or not knowing about these projects)
- Response
 - Retaining and redefining (Camera settings, CDM/LDM, patterns)
 - Learning (DMZ)
 - Publishing (Align > Refine > Design series)

How is NoSQL impacting your jobs?



45

Data Modeling Trend #2

- Knowledge graphs getting the spotlight
- Impacts
 - Data modelers and data scientists collide
 - Traditional modelers extending their skills to include knowledge graphs, or being excluded from exciting machine learning and AI
- Response
 - Fact-based modeling
 - Learning



46

From traditional data modeling to knowledge graphs

- Tables to Resource Description Framework (RDF)
- Attribute value to Triple (Subject-Predicate-Object)
- Surrogate key to URI
- CDM to Ontology
- DDL to OWL
- SQL to SPARQL
- Metadata to "kinds of things"
- Data to "things"
- Benefits: focus on things and make inferences

Pete Stiglich Present July 19, 2023

Anindita Mahapatra Present October 3, 2023

Steve Hoberman Present November 7, 2023

July 19, 2023 Topic "Specialized Database Types and Their Applications"

October 3, 2023 Topic "Data Management and AI/ML"

November 7, 2023 Topic "Top Three Data Modeling Trends"

What topic is Steve Hoberman talking about today?

47

Education: DMZ Sessions

- Unleashing Data's Potential: Mastering Precision with FCOIM. Marco Wobben, CaseTalk
- The Great Escape: Liberating 20+ Years of Legacy Data into Knowledge Graphs and Semantics with AI. Eliud Polanco, Fluree
- Unlocking Enterprise Intelligence: Integrating BI, Bayesian Networks, and OLAP Cubes into Knowledge Graphs. Eugene Asahara, Kyvos
- Avoiding Silos with Semantic Technology Part 1: Introduction to Semantics, Ontology, and OWL. Michael Uschold, PhD, Semantic Arts
- Avoiding Silos with Semantic Technology Part 2: Semantics in Practice. Michael Uschold, PhD, Semantic Arts
- Leveraging data quality checks, predictive modeling, and knowledge graph for credit card fraud detection: A Comprehensive Framework. Rafid Reaz, DMC, Scotia Bank
- Getting Data in and out of a Knowledge Graph. Michael Uschold, PhD, Semantic Arts
- Semantic Knowledge Graphs in Action with Ontotext GraphDB. Sumit Pal, Ontotext
- Zero Copy Integration. Dave McComb, Semantic Arts



48

Data Modeling Trend #2

- More knowledge graphs
- Impacts
 - Data modelers and data scientists collide
 - Traditional modelers extending their skills to include knowledge graphs, or being excluded from exciting machine learning and AI
- Response
 - Fact-based modeling
 - Learning


How are graphs impacting your jobs?



49


Data Modeling Trend #3

- AI experimentations
- Impact
 - Data modelers leveraging AI, or, at some point in the future, AI becoming the data modeler
- Response
 - Learning
 - TLDR AI
 - DMZ



50

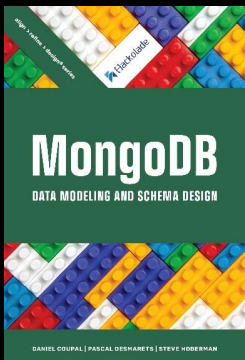
**Create
the initial
model**


 Ok, you are a data modeler tasked with documenting an ideal structure for a database that will be used to manage the operations of an animal shelter. Based on the following user stories, propose an entity relationship design for a single database that will satisfy all the criteria described in the user stories. Include descriptions of the entities, the attributes of each entity including the data type, and a short description of each attribute.

User stories:

1. As a fundraiser for the animal shelter, I want to track all people who have donated to the animal shelter including the amount and date of each donation so that I can plan future fundraising efforts.
2. As a fundraiser for the animal shelter, I want to keep track of each donor's contact preferences including how they prefer to be contacted and whether they have opted out of future communications so that I can effectively and respectfully conduct fundraising outreach.

**What
does AI
have to
do with
Data
Modeling?
Kim
Sever**




 To design an entity-relationship model for a database to manage the operations of an animal shelter based on the provided user stories, we need to identify the key entities, their attributes, and the relationships between them. In this scenario, we primarily need to focus on entities related to donors and their donations. Here's a proposed entity-relationship design:


Entities and Attributes:

1. **Donor Entity:**
 - **DonorID (Primary Key):** Unique identifier for each donor. Data Type: Integer or GUID.
 - **FirstName:** First name of the donor. Data Type: String.
 - **LastName:** Last name of the donor. Data Type: String.
 - **Email:** Email address of the donor. Data Type: String.

51

 Adjust the design to include the following user stories:

1. As a doctor at the animal shelter, I want to keep track of the date when each animal was admitted to the animal shelter, the date and type of each immunization the animal has received, the date of each checkup the animal has had and any treatment notes including medications prescribed, and medical conditions the animal is diagnosed with so that I have a complete picture of the animal's health.
2. As a doctor at the animal shelter I want to ensure that animals are medically cleared for adoption prior to advertising them as available to adopt so that sick animals are not adopted out.
3. As an adoption manager at the animal shelter I want to track the date each animal was admitted along with the date they were adopted.


 To adjust the database design to accommodate the additional user stories related to animal health, adoption management, and volunteer scheduling, we'll need to introduce new entities and expand the existing schema. Below is an extended entity-relationship design that addresses all the user stories:


Entities and Attributes:

1. **Animal Entity:**
 - **AnimalID (Primary Key):** Unique identifier for each animal. Data Type: Integer or GUID.
 - **Name:** Name of the animal. Data Type: String.
 - **AdmissionDate:** Date when the animal was admitted. Data Type: Date/Time.
 - **AvailableForAdoption:** Boolean flag indicating if the animal is available for adoption. Data Type: Boolean.
 - **MedicalCondition:** Text field to describe medical conditions of the animal. Data Type: Text.

**Refactor
the model**

52

 An adopter could be interested in multiple animals and have a different status for each.


 Thank you for the clarification. In that case, we need to adjust the database design to accommodate the fact that an adopter can be interested in multiple animals and have a different status for each. We can achieve this by introducing a many-to-many relationship between the Adopter and Animal entities, along with an additional table to track the status of each adoption interest. Here's the revised design:


Entities and Attributes:

1. **Animal Entity:**
 - **AnimalID (Primary Key):** Unique identifier for each animal. Data Type: Integer or GUID.
 - **Name:** Name of the animal. Data Type: String.
 - **AdmissionDate:** Date when the animal was admitted. Data Type: Date/Time.
 - **AvailableForAdoption:** Boolean flag indicating if the animal is available for adoption.

**Fix
mistakes**

53

 Which of these entities contain content that may be sensitive or subject to regulatory oversight?

 In the context of an animal shelter management database, several entities may contain sensitive or subject-to-regulation data. These entities may be subject to privacy and data protection regulations, especially if they involve personal information about individuals. Here are the entities and specific aspects to consider:

1. **Adopter Entity:** This entity contains personal information about potential adopters, such as their names, email addresses, and phone numbers. Depending on the jurisdiction and regulations, this data may be subject to privacy laws, and the shelter must ensure it is handled securely and with consent from the individuals involved.
2. **Rejected Adopter Entity:** Similar to the Adopter entity, the Rejected Adopter entity also contains personal information about individuals who were rejected as adopters. Privacy

PII consult

54

Education: DMZ Sessions

- What does AI have to do with Data Modeling? Kim Sever
- Data Literacy from an Early Age in the Age of AI: Empower Your Organization. Chandra L. Donelson, Army Intel Deputy Chief Data Officer, Pentagon and Jordan Morrow, DataPrime
- Our Friend the LLM Mini-Hackathon. Kim Sever.
- Mastering Large Language Models: A Practical Workshop. Thomas Niccum, PhD, Iseyon



55

Data Modeling Trend #3

- AI growing, growing, growing
- Impact
 - Data modelers leveraging AI, or, at some point in the future, AI becoming the data modeler
- Response
 - Learning
 - TLDR AI

How is AI impacting your jobs?



56

Other Data Modeling Trends

- Databricks/Snowflake
- Data mesh
- ???



57

Data Modeling Trends

- Trends
 - Mainstream NoSQL projects (“Modernizing”)
 - Knowledge graphs getting the spotlight
 - AI growing, growing, growing
- More connecting and less explanatory
- Sharing over Lecturing
- Your insights as well please!
- Learn at least three new things!



58

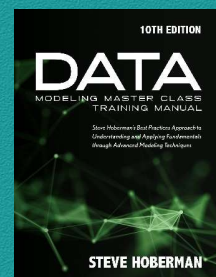
Top Three Data Modeling Trends

Steve Hoberman

me@stevehoberman.com

TechnicsPub.com

DataModelingZone.com



59