



Deelnemerslijst

Business-oriented Data Modelling Masterclass 1 t/m 3 juni 2026

BEDRIJF	NAAM DEELNEMER		FUNCTIE
ABN AMRO	Remi	Mitra	business-analyst
ASN BANK	Mark	van Helvoort	data architect
Banca d'Italia	Antonio Maria	Biasco	business-analyst
Belastingdienst	Robert	Stuffers	other-it-related-job-title
Belastingdienst / Ministerie van Financiën	Martina	Stam	bi-specialist
De Belastingdienst	Jorrit	Janssen	data architect
Deltiq	Dennis	van Gelder	organisation-consultant
FMO (Nederlandse Financierings-Maatschappij voor Ontwikkelingslanden N.V.)	Stephanie	Allegrini	business-analyst
Gemeente Oosterhout	Erik	van Dongen	data engineer
Gemeente Oosterhout	Kheibar	Atai	data engineer
ING	Simone	Giovannozzi	business-analyst
Pensioenfonds Rail & OV	Ties	Martens	data engineer
RVO	Julia	Paranich	information-analyst
RVO Rijksdienst voor Ondernemend Nederland	Nicolette	Put van de	other
Rabobank	Paulina	Bouwmans	other
Rabobank	van	Rijn	general-manager
Rijksdienst voor Ondernemend Nederland	Evelien	van der Grift	other

Welcome Adept Events

www.adeptevents.nl

WHO WE ARE

 **AdeptEvents**

DW & BI SUMMIT

 **BI Platform**

RELEASE

Werner Schoots

Founder Adept Events



BI-Platform.

- Launched in 2008 as online spin-off from Database Magazine (DB/M)
- Topics: Business Intelligence, Data Warehousing, Analytics, Data Management

News

Job board

Selected Whitepapers

Events

Articles

Blogs

Video interviews

Cases


- We welcome your input: redactie@biplatform.nl

  www.biplatform.nl & weekly newsletter

  @BIplatform on X.com & YouTube

  Download the BI-Platform App

 Join our LinkedIn Discussion Group



RELEASE.

- Launched in 1996 as Software Release Magazine
- Topics: Software Engineering – Analysis, Design, Development, Testing and Deployment

News

Job board

Selected Whitepapers

Events

Articles

Blogs

Video interviews

Cases


- We welcome your input: redactie@release.nl

  www.release.nl & weekly newsletter

  @Release_nl on X.com & YouTube

  Download the Release App

 Join our LinkedIn Discussion Group



SEMINARS

All seminars and workshops are organised twice a year

Alec Sharp	Business-oriented Data Modelling Masterclass Working with Business Processes Masterclass Concept Modelling for Business Analysts The Data-Process Connection (virtual half day session)
Rick van der Lans	Ontwerpen van een Nieuwe Data Architectuur
Mathias Vercauteren	Data Governance Sprint AI Governance, Responsible AI and Data Governance – Connecting the Dots
Chris Bradley / Winfried Etzel	Data Management Fundamentals
Lawrence Corr	Agile Data Warehouse Design & Dimensional Modeling
Christian Gijssels	Generatieve AI in Business Analyse
Juha Korpela	Data Mesh – Modeling Data Products and Domains
Thomas Gijssels	AI Agents in de Praktijk – Van Concept tot Implementatie
Multiple speakers	Data Warehousing & BI Summit – Yearly conference in March/April

IN-HOUSE

All seminars and workshops can be organized in-company.
With local speakers and international speakers!



Please contact Werner Schoots

☎ +31 (0)172 742680

✉ seminars@adeptevents.nl

Business-Oriented Data Modelling Masterclass – *Balancing Engagement, Agility, and Complexity*

Presented by
Adept Events and Clariteq Systems Consulting

Alec Sharp
Consultant
Clariteq Systems Consulting Ltd.
West Vancouver, BC, Canada
asharp@clariteq.com
www.clariteq.com


Instructor / course developer background...



Alec Sharp, Clariteq Systems Consulting – asharp@clariteq.com

- 45 years global experience as an independent consultant:

- Business Process Modelling & Business Process Change – discover, scope, analyse, and design/redesign processes
- Application Requirements Specification
- **Data Modelling and Management** My roots!
- Facilitation & Organisational Change
- Project Recovery



- Awarded DAMA's global Professional Achievement Award for contributions to "human-friendly" data modelling

- Author of "Workflow Modeling"
 - best-selling book on process modelling & improvement
 - second edition – a complete re-write

Check out the nice reviews
on Amazon - <http://amzn.to/dHun1o>



Clariteq – small, husband & wife company, global clients

ABB (ASEA Brown Boveri)
Aflac
American Honda
AMP (Australia Mutual Provident)
BackOffice Associates
Bank of Finland
Bellrock
Booking.com
Brisbane City Council (Australia)
Canadian Natural Resources Ltd.
City of Seattle
Civica UK
Clearwater Paper
Corvias
Dell
DHL Express
Dutch National Bank
Elisa
Ericsson
Essity
Eurojust (European Justice Comm.)
European Central Bank
Fortum
Gofore
Helse Vest - Norway
HM Land Registry - UK
Home Depot
Idaho Transportation Dept.
Intel
ISO New England


ING Bank
JP Morgan
Kal Tire
KONE
LGM Financial Services
Liberty Mutual
Livestock Improvement Corp.
MacDonald Dettwiler
Manitoba Public Insurance
Marathon Pipe Line
Microsoft
Ministry of Defence - UK
Ministry of Defence - NL
Ministry of the Interior - Slovakia
MTS Allstream
Nexen
Novo Nordisk
Nusenda Credit Union
OP Bank
Partner Reinsurance
Ritchie Brothers
Phillip Morris
Roche Diagnostics/Pharmaceuticals
Salt River Project
Saudi Aramco
Serco
Shell
Sparta Consulting
State Street Bank
SunGard

SVB (NL)
Synchrotron
Sysdoc
Talent Base
Teck
The MUSIC Group
The Seattle Times
UK Government
University Med Ctr Groningen
YIT(FI)
Washington Gas & Light

– Higher Education –
Carnegie Mellon University
Cornell University
Douglas College
Gonzaga University
Humboldt State University
The Jackson Laboratory
The Ohio State University
Portland State University
Salt Lake Community College
Southern NH University
University of Arkansas
University of British Columbia
University of the Fraser Valley
University of Maryland
University of Utah
University of Washington
Utah Valley University



Background for this course



Overview and logistics

➡	Fundamental and Advanced Topics
1.	Introduction and Level-set - Introduction, Hands-on Case Study - Essentials of Concept Modelling - Transition from Conceptual to Logical, and Logical to Dimensional
2.	Interesting Structures - Types vs. Instances - Recursion, Subtyping, & Generalisation - Meeting New Requirements
3.	Modelling Time, History, & Change
4.	Rules on Relationships and Associations - Multi-way Associatives & Complex Rules - Advanced Normal Forms (4NF & 5NF)
5.	Presentation Techniques for Data Modellers - Core Techniques for Presenting - A Real-life Example

Schedule (CET)

- 09:00 start
- 09:00 - 10:30 class
- 10:30 - 10:45 break
- 10:45 - 12:30 class
- 12:30 - 13:30 lunch
- 13:30 - 15:00 class
- 15:00 - 15:15 break
- 15:15 - 17:00 class
- 17:00 end


Finally...***you***:

- Name (how should I address you?)
- Brief description of your work
- Is there a topic you are especially interested in?
- *Please try to keep your introduction to one minute or less*

What is a Concept Model / Business Object Model / Domain Model...?

- A *description of a business* in terms of
 - **things** it needs to maintain records of – *Entities*
 - **facts about those things** – *Relationships & Attributes*
 - **policies & rules governing those things and facts**
- Models a view of the **real world**, not a technical design (therefore, stable and flexible)
- Can be comprehended by mere mortals (at least initially)
- Graham Witt – “A narrative supported by a graphic”

Graphic component



“Things” first,
data later!

Narrative component

Student definition:

A Student is any person who has been admitted to the University, has accepted, and has registered in a course within a designated time. Faculty and staff members may also be Students.

Plus “Assertions” (policies & rules)


- Each Course is offered through one or more Classes
- Each Class is an offering of a single, specific Course
- Each Instructor teaches one or more Classes
- Each Class is taught by one Instructor (which may or may not be true...)

Many rules can't be shown on the diagram...

- A Student can not register in two Classes of the same Course in the same Academic Term

A better looking version of the model on the previous slide

Independent entities across the top



Case study – Concept Model, Services, Use Cases, Business Processes

Client –

- Regulatory agency ensuring the safe design, installation, and use of technical equipment
- Natural gas systems, electrical systems, boilers and pressure vessels, elevating devices, & many more




Goal –

- Shift from an inspection-based model (~800 inspectors!) to client-managed safety programs
- Clients will apply for a *Client Safety Management Program Authorisation (CSMP Authorisation)*
 - must show effective processes and accurate record-keeping
- Clients will pay a fee for managing *their own safety programs!* Still beneficial!



Case study – Concept Model, Services, Use Cases

- Business Development chooses Pilot Program – boilers and pressure vessels in Oil & Gas fields



- Current systems won't support CSMP, time-consuming and expensive to change them – IT and Finance suggest 18 – 24 months of work
- BD is unimpressed by IT and Finance objections (“You're being mindlessly obstructionist!”) and proposes work-around procedure. *Guess which tool they intend to use?*
- I'm hired to identify end-to-end implications – “Design a process and determine IT requirements that will allow this procedure to work.”
- *Concept Modelling was a critical tool in understanding the underlying policies, and developing the process & requirements*

Building your initial Concept Model, step-by-step

Identify and define "Things"

1. Collect terms

- 1:1 interviews
- survey (e.g., email)
- group brainstorm
- analyse documents

2. Isolate "things"

Ask *Is this...*

- a thing?
- a fact about a thing?
- or "other stuff?"

3. Identify synonyms

- select a term to use
- as general as possible
- just for this initiative, not the entire enterprise

4. Define each thing

- "good enough for now"
- first, identify "anomalies, sources of confusion, and valid differences of opinion"
- select which to include

Develop initial Concept Model

5. Organise things

- independent things across the top
- then laid out top-down by dependency

6. Draw relationships

- show dependency
- parent-child drawn bottom-to-top
- otherwise, side-to-side

7. Name relationships

- in both directions
- active verb-based!
- not mushy – *has*
- not meaningless – *related to*

8. Add cardinality

- use words first
- 1:1 is probably wrong
- 1:M (one to many)
- M:M (many to many)

Refine Concept Model

9. State assertions

- forcefully, for each relationship
- challenge the assertions!
- restate the assertion & why it changed, if it did

10. Redraw the model

- shows revised assertions
- e.g., 1:M becoming M:M
- e.g., dependent things becoming independent

11. Collect attributes

- a few for each thing
- not *all* attributes
- don't worry about normalisation

12. Move to identifying:

1. events / services
2. use cases / user stories

"Terminology Analysis" – usually a great place to start

From one-on-one interviews with 8-10 key stakeholders we gathered ~200 terms related to CSMP (Client Safety Management Program) – "anything that went by a name."
Here are 24 that met the criteria to be a "thing" – the candidate *Business Objects*.


Device	Client	Unit	Location	Company	Site
Applicant	Pressure Vessel	Operator	Owner	Boiler	Licensee
Slug	Operation	Verification	Customer	Plant	Inspection
Pig	Facility	Permission	Authorisation	License	Confirmation

To isolate the "things"
(Business Objects) we went
through the 200 terms, each on a
Post-it, and asked "Is this a...

- ...thing?"
 - ...a fact about a thing?"
 - ...or other stuff?"
- (forms, systems, job titles,
spreadsheets, activities,
organisations, etc.)

3 – Identify synonyms and select a term to use for the initiative

Let's arrange these terms into columns of synonyms. It's always a surprise for the business to see how many terms are used to describe the same fundamental thing!



4 – Agree on a "good enough" definition" and important constraints

Agreed terms
("core nouns" or
business objects)



Definition:


A CSMP Authorisation is a permission (or license) to operate a self-managed safety program (a Client Safety Management Program) at a specific Facility for a specified time period, usually 1, 2, or 5 years.

Constraint (Assertion):


The CSMP Authorisation is "all or nothing" – it covers all the Units at a Facility.

Business Object Model Version 1; not perfect, but a good start


5 – organise objects top-down by dependency




6 – add relationship lines indicating dependency



7 – name relationships in each direction (verb-based!)




8 – only then add cardinality (1:M, M:M)



Step by step!

Just boxes and lines, but raises important questions



What do we issue the Authorisation to?


Are Units permanently part of one Facility?

What do we Inspect?

9 – State each relationship forcefully as an assertion, and challenge it


Now, state the relationships **emphatically** as Assertions. **Each** Client operates **one or more** Facilities! Then, **challenge** them! Again, don't worry yet about **optionality** – whether the relationship **must be** or **may be** be present.

We only care now about the **maximum** – each ObjectA is related to a **maximum of one or one or more (or many)** ObjectB.




9 – Revised assertions from challenges

Now, state the relationships **emphatically** as Assertions. **Each** Client operates **one or more** Facilities! Then, **challenge** them! Again, don't worry yet about **optionality** – whether the relationship **must be** or **may be** be present. We only care now about the **maximum** – each ObjectA is related to a **maximum of one or one or more (or many)** ObjectB.



10 – redraw the Business Object Model based on revised assertions

Now we will re-draw the initial Concept Model based on changes that came from challenging the Assertions in Ver. 1.




Note:

You don't always get what you *want* or what you think is the *right* thing in Concept Modelling. In this case the client (the Regulator) said they always wanted a Facility to be operated by ONE AND ONLY ONE Client. If a Facility was operated by multiple Clients, they would require the Clients to form a new Joint Venture Client. This was to ensure that if there were legal difficulties, there was only ONE Client to go after. Or, as they put it, "one throat to choke."

Later in the project, they realised they needed a history of the Clients that had operated a Facility, so the Client-Facility relationship became Many-to-Many, and Facility was modelled (correctly) as an independent Entity, as shown here:



11 – Collect some (not all!) attributes



Sketching this out was *fast*, and raised important questions that had not occurred to the client!

Highlights:

Model took
~90 minutes


This was done initially without
any data modelling terminology
or symbols!

It's not perfect, but the
businesspeople found it
extremely useful.

Don't worry now about "normalisation" at this point –
that is done in Logical Data Modelling.
Then, "Unit" will become at least 5, more granular entities –
Unit, Manufacturer, Manufacturer Make/Model,
Unit-Facility Registration, and Risk Factor.

12 – Identify Services (Events) then Use Cases / User Stories

Finally, we identified the necessary Services (verb-noun pairs) and Use Cases / User Stories that would access the Services. We developed high-level ("upper conceptual") descriptions to provide the vendor with core requirements and avoid a bulleted list requirements document.



Note – "User Story" and "Use Case" are not so different

Different format and detail, but the same basic concept.
Initially, at the Scope level, they're much the same:

User Story (who – what – *why*):

"As a Client, I need the ability to Register Unit(s)
so I can maintain compliance with my CSMP Authorisation"


Use Case: (who – what – *how*):

"Client Register Unit via Portal"

When we add detail at the Concept level, they become identical:


- User Story / Use Case abstract
- Main success sequence – dialogue in "when-then" format
- Alternate sequences – variations, exceptions, errors

Clarify scope of the new process and identify participants




Process Summary Chart – simplified “what,” plus “who”

The initial, business-friendly workflow model




A “Handoff Level”
workflow model




Mission accomplished! Conclusions:

- "Plan A" rejected – agreement that Unit data *must* get into S-MAN
- "Plan B" (change the app) looks good, but the vendor estimates are *HIGH*
- "Plan B Minus" (existing functionality plus CSR work) is *worth the cost*



1. If requirements, issues, assumptions, etc. are in lists, people will argue endlessly; if they are in an *integrated* and *understandable* set of models, it's much harder to dismiss the reality of the situation
2. Process Models, Use Cases, Service Specs, & Concept Models: **essential!**

Data Models in an integrated, model-based framework




Everything relies on the Concept Model / Data Model

All use the language and constraints of the Concept Model (the “thing model”) – the ultimate “what”


The university is initiating the “Strategic Enrollment” program to raise Student graduation rates in part by ensuring Classes are available for Student registration when needed.

Business Goals
& Objectives


Business
Process




Presentation
Layer
(user interface)



Application
Layer
(rules & logic)



Data
Layer
(data & storage)



Use Cases/User Stories:
- Who (Actors) needs access to the Services, and how (Platform)?

Use Case
actor + service + platform:
Advisor Register Student in Class via SRS

Verb-Noun pairs:

- The Services (event-handlers) that are at the heart of a Service Oriented Architecture.
- Also "building blocks" of Business Processes


Service
verb + noun (+ noun):
Register Student in Class

Entity
noun:
Class

The core Nouns or Things in your enterprise. Also known as Business Objects.

Bonus – great starting point to discover your Events/Services and Use Cases/User Stories


With progressive detail, Concept Modelling supports Agile



Techniques and methodologies

- The same techniques are used in different sequences, with different emphasis, in different situations.
- Concept Modelling to clarify language is a great starting point.


Larger project (enterprise-wide, cross-functional): process-oriented / “outside-in” –



These are typical overall flows:

- there are many variations
- there is always much iteration

Smaller project (local, departmental): service or use case-oriented / “inside-out” –



A blank slide to maintain balance and page numbering

Overview and logistics

➡	<i>Fundamental and Advanced Topics</i>
1.	Introduction and Level-set • Issues, Principles, Hands-on Case Study • Essentials of Concept Modelling • Transition from Conceptual to Logical, and Logical to Dimensional
2.	Interesting Structures • Types vs. Instances • Recursion, Subtyping, & Generalisation • Meeting New Requirements
3.	Modelling Time, History, & Change
4.	Rules on Relationships and Associations • Multi-way Associatives & Complex Rules • Advanced Normal Forms (4NF & 5NF)
5.	Presentation Techniques for Data Modellers • Core Techniques for Presenting • A Real-life Example

The basics: ERA – Entities


An *entity* is a distinct thing the business *needs* to know about - a *person*, *place*, *thing*, *event*, *concept*, or *organisation*, and...

- ☆ is named with a *singular noun* that implies a single instance
 - not a plural or collective noun, list, set, collection, report, etc.
 - we can discuss “one of them,” e.g. “Weather” is not a good name
- has multiple occurrences (or instances)
- we *need* to and *can* keep track of (differentiate) each occurrence
- has *facts* that must be recorded, e.g.
 - *Student* attributes: Number, Name, Birth Date, Major, GPA, ...
 - *Student* relationships: “majors in” *Subject*, “enrolls in” *Section*
- is acted on by *processes*, so they make sense in a “verb-noun” pair
- ☆ refers to the *essence*, not the *implementation* – *the most common error is to identify artifacts (forms, reports, spreadsheets, ...) as entities!*

Named - a business-oriented noun / noun phrase

Defined - “What is one of these things?” or
“What do you mean by _____?”

☆ These are the ones our business partners often struggle with.




Naming and definition – the essence of Concept Modelling

Agreement on naming is essential – entities are the *nouns* at the heart of business language. They are what processes act on, applications manipulate, databases record, and BI & analytics tools provide info about.

So, organisations need a *common language* more than ever, for...


- Data integration (data lake, data mesh, data fabric, data virtualisation, data warehouse, ...)
- SLMs (*Small* Language Models vs. *Large*)
- Mergers/acquisitions/partnerships/...
- Business analysis – most requirements can't be stated without using a term from the Concept Model
- Performance measures, e.g., KPIs

Note – it works best if you don't start by talking about *Concept Modelling* or *Data Modelling*...




“Now! That should clear up
a few things around here!”

Concept Modelling principles



“Let's start here with
Special Tax Rate Variation Comment Type...”



Models should:

- Mask unnecessary detail
- Highlight what matters
- Use visual cues consistently

We will focus on:


- Directionality
(top-down by dependency)
- Simplicity and abstraction
- Minimizing graphic "widgets"

The basics – ERA – Relationships

An association between Entities that the business must keep track of

Named in both directions

- verb-based phrase
- the line tells us they *are* related,
- the name tells us *how*




Different types of relationships

1. parent-child or characterising – “bottom to top” relationship from an entity to a dependent entity (1:M)
2. associating – “side to side” relationship between entities that are not dependent on one another (usually M:M)
3. classifying – “side to side” relationship from reference data to the classified entity (seldom shown in the Concept Model)

Dependency is shown top down – No Dead Crows


Relationships have rules

- cardinality – 1:1 (almost certainly wrong,) 1:M, M:M
- optionality – relationship *may* be present or *must* be present (not shown until later, in the logical model)



Relationship cardinality (maximum cardinality)

A kind of “business rule”
that applies to relationships




To determine cardinality, *first name the relationships properly, and only then:*

- for each entity, ask
“Can one of these be related to a *maximum* of One of the other or a *maximum* of *Many* of the other?”
- record the answer (One or Many) at the “other” end;
“One or More” works better for businesspersons than “Many”
- possibilities – 1:1 (error), 1:M (common), M:M (more work, eventually)

Relationships – state as assertions

1. You *must* state the relationship name as an assertion, in both directions (for clarity and confirmation)
2. Be clear on whether cardinality is “one” or “one or more” (don’t worry about “may” and “must” at first)
3. *Emphatically* begin the assertion with the word “Each”
4. Try it on this model...



Note –

A Class is a scheduled offering of a Course during an Academic Time Period, e.g. a Semester or an Academic Year.
During an Academic Time Period there may be one or more Classes for a Course.
Each Class is held on specific Days (e.g. Monday & Wednesday,) at specific Times (e.g. 10:30-11:30,) in specific Rooms (e.g. AQ3100 & CC7232.)

Each Instructor teaches one or more Classes
(Sounds good...)


Each Class is taught by one Instructor...

1. Student-Class
2. Course-Class
3. Instructor-Class
4. Room-Class

Which ones might be *incorrect*?

Discussion – state as assertions, identify incorrect ones

In some universities, Students in the same Class could be earning credit for *different* Courses – it could be a M:M relationship.



Each Class is taught by One or More Instructors. On what basis?

- team teaching
- backup
- replacement
- specialist
- guest lecturer
- lab assistant
- teaching assistant
- ...

We are discovering reference data to describe an Instructor's Role.

All of this has an impact on the Business Process! It's easier to resolve these rules before working on the Process.

1. Student-Class
Each Student registers in one or more Classes ✓
Each Class is registered by one or more Students ✓
2. Course-Class
Each Course is offered via one or more Classes
Each Class is an offering of one Course ? – depends on Policy
3. Instructor-Class
Each Instructor teaches one or more Classes
Each Class is taught by ~~one~~ One or More Instructors
4. Room-Class
Each Room is the location of one or more Classes
Each Class is located in ~~one~~ One or More Rooms

The basics: ERA – Attributes

A fact about an entity recorded as a piece of data.

If facts are needed about a relationship, we will later (in the Logical Data Model) create an entity that represents the relationship and records its facts

Like Entities, attributes are named and defined

Not every possible fact – just the ones we need

Have properties that we address during the transition from Concept Model to Logical Data Model

1. base or fundamental attribute
2. single-valued vs. multivalued – one attribute can have multiple values, at a time or over time
3. fundamental vs. redundant – the same value is recorded multiple times in different entities
4. “user-entered” vs. constrained – attribute can only come from a limited set, as in a drop-down list


Traditionally alphanumeric data; now includes richer types e.g., retinal scan image or voice audio clip

Eventually, in the logical model, an entity will contain only base / fundamental / *essential* attributes:

- an *essential fact* about that thing (entity)
- *not* multi-valued
- *not* redundant

(a redundant attribute is an attribute that is really an essential fact about a *different* entity, so its value is recorded multiple times, redundantly)

- and *not* derived or calculated from other attributes; otherwise, clearly flagged “derived”



E.g., Mobile Number,
Office Telephone Number,
Facebook ID,
LinkedIn ID, ...

Summary – three types of data models

A more detailed
version is on Page 74

Different levels of detail support different perspectives

1 Contextual (Scope)	2 Conceptual (Overview)	3 Logical (Detail)
✓ <i>Context model</i> ✓ Agreement on “big picture,” context, and some vocabulary ✓ A block diagram of “subject areas,” higher level than individual entities ✓ Shows the scope or “footprint” ✓ Optional – not useful on smaller projects	✓ <i>Concept Model</i> ✓ Agreements on basic concepts, vocabulary, and rules	✓ <i>Logical Data Model</i> ✓ Complete detail for physical design


Some important differences

✓ Main (“recognisable”) entities only - a singular noun used daily ✓ Main attributes only, many are non-atomic ✓ M:M relationships ✓ Doesn't show keys ✓ Not normalised ✓ A “one-pager”	✓ All granular entities – many too detailed to come up daily ✓ All attributes included, all are atomic ✓ All M:M resolved ✓ Shows primary & foreign keys ✓ Fully normalised ✓ Five times as many entities
---	---


My most plagiarised
diagram ever!

For reference – the Information Engineering symbol set


- This symbol set was refined and developed by Clive Finkelstein.
- Known in some tools as the "Martin IE" symbol set.
- Strengths are:
 - symbols are not "overloaded" – they explicitly convey only *one* idea.
 - can show as much or as little as needed in terms of rules.



A natural progression



Different ways to get started



Some advice on starting the concept model



Don't begin with a lecture
on data modelling
(but I have a painful story
that had a happy ending)

If you can, don't
even mention "data
modelling"

We use "terminology analysis" –
starting with the nouns –
at the outset of every project.
This was demonstrated earlier in the
Client Safety Management example.

Starting a Concept Model bottom-up

- 1) Interview business representatives about their area:
mandate and activities, goals and objectives, issues and opportunities,
needs and wants, likes and dislikes, etc....
Nod sympathetically, but ignore it all (almost!)
Instead, capture “terms” – anything that goes by a name.
- 2) Later, write each term on a large Post-it
- 3) In a facilitated session, participants sort terms into categories:
 - Things (entities, but don’t use the term... yet)
 - Facts about things (add new “thing” if it's not there already)
 - “Other stuff”

As needed, introduce criteria to
be a “thing” (an entity)

“Other stuff” includes:

- Metrics
- Organisations, departments, jobs, roles, ...
- Processes, functions, activities, tasks, ...
- Systems, tools, equipment, mechanisms, ...
- Reports, forms, screens, queries, ...
- Other – too vague, only one instance,
a “fact of life,” not a thing we track, etc.

Exercise 1: building an initial Concept Model

The assignment:

The following describes project tracking at Amalgamated Automaton. Read it over and be prepared to discuss the things about which the business needs to record information, and the important facts about them. The instructor will lead the development of an initial data model.

Amalgamated Automaton, Inc. has a growing Information Systems department. Until recent years, the department was concerned almost entirely with selecting, installing and maintaining purchased software packages. Recently, however, the focus has shifted towards the in-house development of application software.

One of the problems confronting the IS department is that they have no base of historical data to aid in trend analysis or estimating development effort, nor any effective means of charging back development costs. The proposed solution is to develop a simple Project Tracking System, which will work in conjunction with the existing Personnel and General Ledger Systems.




When a development project is initiated, a project name and a short description are recorded, among other things. Soon, before any further work is done on the project, a new account is created on the G/L System, identified by a G/L account number. Project costs will be charged to this account, and the project budget is recorded as the initial account balance in dollars.

Project planners break a project down into many tasks, perhaps hundreds. A typical project task might be "Test Order Entry Module". Some of the facts which are required about tasks include a brief task description, estimated work hours, and the scheduled start and finish dates.

Eventually, individual employees are assigned responsibility for the tasks. Some tasks will be the responsibility of many employees, and an employee might be assigned to many tasks. As each employee is assigned to a project task, their planned start and finish dates, their contribution to the task (not a "kind of work," but their specific duties on the task – e.g., "Develop test scripts"), and the estimated number of hours they are to spend on the task are recorded. Employee information such as the employee name and number are available from the existing Personnel System, although it will have to be modified to record the employee's hourly charge out rate.

When an IS employee begins work on a new task, their actual start date is recorded. A running total of the number of hours that they have worked on each started task is updated regularly. At the same time, the remaining balance in the project account is updated. When an employee completes a task assignment, the actual completion date is recorded.

Workshop example



Introduce "thing criteria" as necessary:

- *singular noun* – can talk about one of them (Worker not Staff, Item not Inventory)
- *multiple instances*
- must need to and be able to track each instance (uniquely identify each)
- has facts that must be recorded
- makes sense in a "verb-noun" pair
- NOT an artifact like a spreadsheet or report (not a Call Log or Worker Directory or...)

Identifying Entities – three common errors

1. Treating an “artifact” (a spreadsheet, report, web page, form, etc.) as an Entity – an Entity is a fundamental, singular thing with no reference to *implementation* – Artifacts reflect implementation (form, DB, spreadsheet, list, ...) and contain attributes from
 - *multiple different Entities* or
 - *multiple instances of the same Entity*e.g., “Admission Request Form” or “Orders Summary Spreadsheet” or “Daily Call Log” or “Materials List Fax” or “Class Roster” or “Course List” or...
2. Identifying an Entity that exists in the real world, but whose *instances* can't be uniquely identified e.g., “*Transit System Passenger*” or “*Event Attendee*”
3. The “types vs. instances” problem – failing to clarify if the Entity deals with *types* of things (or *categories* or *kinds* or *classes* of things) vs. specific *instances* of things e.g., “*Test*” – is this a *type* of Test, or a *specific instance* of a Test? *more examples coming...*

Types vs. Instances example

Specifications for a
type of vehicle –
a 2011
Volkswagen GTI



5dr HB DSG

\$31,275 MSRP

✓ Gasoline

✓ 6.3L - 8.7L fuel / 100 km

✓ Front-Wheel Drive

✓ 200 horsepower

Mechanical	
Engine	Turbocharged Gas I4
Displacement	2.0L/121
Fuel System	Electronic Fuel Injection
Horsepower	200 hp @ 5100-6000 rpm
Torque	207 ft-lb @ 1800-5000 rpm
Steering	Power Steering
Drivetrain	Front-Wheel Drive
Transmission	6-speed DSG transmission

Interior	
Passenger Capacity	5
Front Head Room	998 mm
Front Leg Room	1,046 mm
Front Shoulder Room	1,384 mm
Rear Head Room	978 mm
Rear Leg Room	902 mm
Rear Shoulder Room	1,361 mm

Failing to distinguish these is a
common modelling error.



Alec's 2011 Volkswagen GTI
An "instance" entity – Vehicle

A "type" entity – Vehicle Make/Model

Types vs. Instances – “What do you mean by a Bus?”



A category of Bus – a "meta-Type?"
(transit, articulated, intercity, minibus, ...)
A Make and Model of Bus – a Type?
An individual Vehicle? – an Instance?

Model	Length	Width	Introduced
Xcelstor ^[18]	35 feet (11 m)	102 inches (2.6 m)	2008
	40 feet (12 m)		
	60 feet (18 m)		
MIDI	30 feet (9.1 m)	96 inches (2.4 m)	2013
	35 feet (11 m)		

“What do you mean by a Bus?”

254 British Properties



Inbound From Glenmore and Bonnymuir via Bonnymuir, Stevens, Taylor Way to Park Royal terminus (extends to Downtown Vancouver during Monday-Friday peak hours).

Outbound From Park Royal (from Downtown Vancouver during Monday-Friday peak hours) via Marine Drive, Park Royal South, Taylor Way, Southborough, Eyremount, Cross Creek, Chartwell, Crestwell, Eyremount, Fairmile, Southborough, King Georges Way, Robin Hood, Kenwood, St. Andrews, Bonnymuir to Glenmore terminus.

Park Royal to British Properties and return to Park Royal

MONDAY TO FRIDAY																	
Vancouver		Park Royal		Leave Eyremount at Highland		Leave Bonnymuir at Glenmore		Leave Eyremount at Highland		Leave Marine at 14th		Arrive Park Royal		Arrive Downtown Vancouver		Connecting Buses	
6.35	6.53R	6.53R	7.03	7.15	7.31	7.34	7.34	7.54	7.54	7.54	7.54	7.54	7.54	7.54	7.54	7.54	
6.45	7.23R	7.23R	7.33	7.45	8.01	8.04	8.04	8.24	8.24	8.24	8.24	8.24	8.24	8.24	8.24	8.24	
7.47	8.07R	8.07R	8.17	8.28	8.44*	8.47	8.47	9.16	9.16	9.16	9.16	9.16	9.16	9.16	9.16	9.16	
8.20	8.40	8.53	9.06	-	-	9.15P*	9.15P*	9.41	9.41	9.41	9.41	9.41	9.41	9.41	9.41	9.41	
9.22	9.47P	10.00	10.13	-	-	10.22P*	10.22P*	10.43	10.43	10.43	10.43	10.43	10.43	10.43	10.43	10.43	

254 British Properties



- A Bus Route?
- A Bus Route Scheduled Departure
- An instance of a Bus Route Scheduled Departure?

Never be afraid to ask “What do you mean by....?”



Discussion – good *Entity* or *not*?

Which of the following might ***not*** be valid entities?
And if not, *why* not?

Transcript

Student

Building

Student
Directory

Faculty
Member

Instructor
History

Department

Course

Organisation
Chart

Prerequisite
List

Payment

Student
Body

Class
Roster

Scholarship

Faculty

Assistant
Dean

Admission
Date

Phillips
Building

Registration

Section

Course
Catalogue

Physics

Class

Professor

Admission
Request
Form

And a bonus...

Time

Discussion – good Entity or not?

Which of the following might **not** be valid entities?
And if not, *why* not?

☒ Transcript	☒ Student	☒ Building	☒ Student Directory	☒ Faculty Member	☒ Instructor History
a report			a report		a list, "history" is not singular, and a history of <i>what</i> ?
☒ Department	☒ Course	☒ Organisation Chart	☒ Prerequisite List	☒ Payment	☒ Student Body
		a visual report	a list		not singular
☒ Class Roster	☒ Scholarship	☒ Faculty	☒ Assistant Dean	☒ Admission Date	☒ Phillips Building
a report			a Job Title	an attribute	an instance
☒ Registration	☒ Section	☒ Course Catalogue	☒ Physics	☒ Class	☒ Professor
		a report	an instance		a Job Title
☒ Admission Request Form			☒ Time		
a form (artifact)					not a distinct thing ("what is a <i>time</i> ?" but can you think of any time-related entities?

Entity definition – bad example then a good format

~~Customer~~

~~We have a variety of Customers that operate in multiple geographies, and these must be tracked in order to consolidate purchasing statistics and enable our rating process to identify our best Customers.~~

Not a good definition

- Interesting background and miscellaneous points
- *Doesn't* answer the question “What is *one* of these things?”

Customer

1. A Customer is a person or organisation that is a past, present, or potential user of our products or services.
2. Current examples include Soletron (contract manufacturer,) Cisco Systems (OEM,) Arrow Electronics (distributor,) Best Buy (retailer,) M&P PCs (assembler,) and individual consumers.
3. Excludes the company itself when we use our own products or services but includes cases where the Customer doesn't have to pay (e.g., a charity.)

Entity definition format:

1. A description of which real-world things will be included in scope.
This might be developed from a list of standard “thing types” – person, organisation, request, transfer, item, location, activity, etc.
Be sure to identify any specific inclusions (“This includes...” or “This is...”)
2. Illustrate with examples:
 - 5 – 10 sample instances
 - diagrams or scenarios
 - illustrations such as reports or forms
3. Interesting points – anomalies, synonyms, common points of confusion, etc.
May include specific exclusions (“This excludes...” or “This is not...”)

Entity definition - start by opening minds

“Can anyone think of examples that might surprise someone else – that is, anomalies or potential sources of confusion.”

E.g., how could we legitimately have different ideas what “Employee” means?

Employee

Project

Account

Task

Starting an Entity definition

“Can anyone think of examples that might surprise someone else – that is, anomalies or potential sources of confusion.”
E.g., how could we legitimately have different ideas what “Employee” means?

F/T vs. P/T?

Only IS Department?

Include management, or only individual contributors?

Still in recruitment (an applicant)?

Onboarded? on probation? active? retirees?

Include contractors, student interns, vendor staff, etc.?

Volunteers?

A type of worker (DBA or tester) or a specific person?

A robotic, automated, or AI agent?

Employee

Project

Account

Task

Starting an Entity definition

“Can anyone think of examples that might surprise someone else – that is, anomalies or potential sources of confusion.”
E.g., how could we legitimately have different ideas what “Employee” means?

Let's see what the group decided

- F/T vs. P/T? – Both
- Only IS Department? – No
- Include management, or only individual contributors? – Yes, everyone
- Still in recruitment (an applicant)? – No
- Onboarded? on probation? active? retirees? – Yes, all
- Include contractors, student interns, vendor staff, etc.? – Yes, all
- Volunteers? – Yes
- A type of worker (DBA or tester) or a specific person? – No, only a specific person
- A robotic, automated, or AI agent? – No, only a real person

Employee

Project

Account

Task

Defining the Entity "~~Employee~~" – "Worker"

Definition format:

1. A description of which real-world things are within in scope, and any specific inclusions ("This *includes...*" or "This *is...*")

2. Illustrate with examples – 5 to 10 sample instances or types

3. Interesting points – anomalies, synonyms, common points of confusion, etc.
May include specific exclusions ("This *excludes...*" or "This *is not...*")

Worker (renamed from Employee):

A *Worker* is a person, whether or not directly employed by *the company*, but with some sort of employment contract or arrangement, who has been or may be assigned to a Project.

Worker includes:

- Full or Part-time Employees who have been onboarded, including Probation, Active, Seconded, Suspended, Retired....
- Contractors
- Consultants
- Student Interns
- Vendor Staff Persons
- Company Owners and Managers

Key points:

- "Worker" was chosen as the entity name because it is more generalised than "Employee."
- A Worker may not necessarily be billable on a Project, e.g., a non-chargeable Subject Matter Expert or Volunteer
- Worker excludes:
 - Job Roles, e.g., DBA or Technical Writer
 - Robotic, Automated, or AI Agents (this might change)

Another example – starting an entity definition for Task

“Can anyone think of examples that might surprise someone else – that is, anomalies or potential sources of confusion.”
E.g., how could we legitimately have different ideas what “Task” means?

-
-
-
-
-

Worker

Project

Account

Task

Another example – starting an entity definition for Task

“Can anyone think of examples that might surprise someone else – that is, anomalies or potential sources of confusion.”

E.g., how could we legitimately have different ideas what “Task” means?

Key points that typically arise:

- A *type* of Task or a specific Task?
(the types vs. instances problem)
- Part of a specific Project or
used across multiple Projects?
- Produces a specific deliverable or state?
- Time-bounded or ongoing?
- Performed by one Worker or
one or more Workers?
- ...

A **Task** is a specific, time-bounded, unit of work, within a single Project, intended to be performed by one or more Workers, that produces an intended deliverable or achieves a specific state.

Examples:

- Code *Place Order* service
- Test *Place Order* service

Excludes:

- types of Tasks
- ongoing (non time-bounded) activities such as management or administration


Worker

Project

Account

Task

Now we have definitions – it's "safe" to draw the ER model



First arrange entities top-down by dependency.
Then add relationships with a verb-based phrase.
Then add cardinality (1:1, 1:M, M:M.)

The diagram illustrates the relationships between three entities: Course, Class, and Instructor.


- Course Entity:** Attributes include Course Code (Department), Course Number, Course Description, Department Name, Department Code, Department Tel. Number, Faculty Name, Credit Hours, and Prerequisite Courses.
- Class Entity:** Attributes include Course Code (Department), Course Number, Course Description, Credit Hours, Prerequisites, Class Number, Delivery Method, Maximum Enrollment, Instructor Name, Lectures (Days, Times (start - end), Room Numbers), and Tutorials (Days, Times (start - end), Room Numbers).
- Instructor Entity:** Attributes include Instructor ID, Name, Level, Section Role, Section Committed Weekly Hours, and Annual Evaluation Rank.

Relationships:

- Course to Class:** "is an offering of" (1:M relationship).
- Course to Instructor:** "is assigned to" (1:M relationship).
- Class to Instructor:** "is assignment of" (1:M relationship).

Relationships – a few more points

A significant, named association between entities – one of the types of facts about entities that data models depict




Guidelines


- named with a descriptive, verb-based phrase – not "has" or "is related to" (the line tells us they *are* related; the name tells us *how*)
- named in both directions – try to use the same root word at both ends (e.g., "classifies" and "is classified by")
- the complete name reads like a sentence (noun verb noun) – "Position is classified by Job Category"

1:1 relationships – *almost always an error!*

- Note – a 1:1 relationship might be necessary in the Physical Database Design e.g., “Fixed Asset” records financial data about a “Network Component” but they are in two separate systems (the G/L System and the Configuration Management System) connected by a 1:1 relationship





- ✗ Incorrect analysis
e.g., Project costs are probably prorated across *many* Accounts



- ✗ Failing to account for changes over time
e.g., an Employee may hold only *one* Credit Card at a time, but *many over time*, and we virtually always want history.
The most common written constraint in Concept Modelling is “*one at a time but many over time*.”






Relationship optionality (logical models only)

Each Building is the location of a *minimum* of Zero Positions

(and a *maximum* of Many Positions)

Each Position is located at a *minimum* of One Building

(and a *maximum* of One Building)



or,

Each Building *May Be* the location of *One or More* Positions

Each Position *Must Be located at* *One* Building

To determine optionality (a.k.a. minimum cardinality)

- for each entity ask "Can one of these be related to a *minimum* of Zero or a *minimum* of One of the other entity?"
- record the answer – 0 or 1 – at the "other" end
- "zero" means an optional relationship (*May Be*) and
- "one" means a mandatory relationship (*Must Be*)
- easier form: "Each one of these *May Be* be or *Must Be* related to the other?"

Don't forget the four Ds of Data Modelling

1

Definition

- “What *is* one of these things?”
- List common and unusual instances
- “Are there any known anomalies?”
- “What are the potential differences of opinion?”

2

Dependency

- “What type of entity is this?”
- “What other entity does it depend on?”
- Essentially
 - is it a free-standing thing?,
 - is it a type of thing?,
 - is it repeating detail about some other thing?

_____ Please let us know the key point (or points) _____

that mattered most to you in this first section.

3

Detail

- Don't dive into detail – keep it in its place!
- GEFNI* HPDL!**

**Good enough for now!*

***Hard part, do later!*

4


Demonstrate the Data

- Assertions / narrative rules
- Sample data values or instances
- Scenarios or use cases
- Props (e.g., report layouts or common documents)
- Also demonstrate the *value* of the model

Overview and logistics

➡	<i>Fundamental and Advanced Topics</i>
1.	Introduction and Level-set • Issues, Principles, Hands-on Case Study • Essentials of Concept Modelling • Transition from Conceptual to Logical, and Logical to Dimensional
2.	Interesting Structures • Types vs. Instances • Recursion, Subtyping, & Generalisation • Meeting New Requirements
3.	Modelling Time, History, & Change
4.	Rules on Relationships and Associations • Multi-way Associatives & Complex Rules • Advanced Normal Forms (4NF & 5NF)
5.	Presentation Techniques for Data Modellers • Core Techniques for Presenting • A Real-life Example

Phase 2 of three phases in data modelling




From conceptual to initial logical

The progression from conceptual to logical is largely based on identifying and dealing with three attribute characteristics

- Multi-valued - the attribute can have multiple different values for one instance of the entity, either “at a time” or “over time”
E.g., “Employee Name” if aliases or previous names are tracked
 - move it **down** to the “many” end of a 1:M relationship into a characteristic entity
 - if it's a fact about a M:M relationship between entities, move it down to the “many” end of a 1:M relationship into an associative entity
 - this puts the data structure into 1st Normal Form – 1NF
- Redundant - the same attribute value is recorded multiple times, in different entity instances, possibly inconsistently
E.g., “Company Name” in a “Department” entity
 - move it **up** to the “one” end of a M:1 relationship to one of the parent (or higher) entities (2nd Normal Form – 2NF)
 - You might have to create a new parent entity where none existed before
- Constrained - a descriptive attribute needs to be restricted to a set of standard (or “allowable”) values to improve integrity and reporting
E.g., “Employee Type”
 - move it **out** to the “one” end of a M:1 relationship to a reference or other related entity (3rd Normal Form - 3NF)

A simple Concept Model to Logical Data Model example



In this example we:


- resolve the M:M relationship between Instructor and Class
- move redundant Department attributes in Course up into a new Department entity
- create a reference entity to standardise the values of "Assignment Role"

ER diagram illustrating a university database schema:

- Entities and Attributes:**
 - Student:** Student ID, Name, etc.
 - Course:** Course Code, Course Number, Course Description, Department Name.
 - Class (or Section):** Class Number, Max. Registrations, Days / Times, Locations.
 - Registration:** Registration Date, Registration Status, Midterm Grade, Final Grade.
- Relationships:**
 - registers in:** Connects Student and Class (or Section).
 - offered via:** Connects Course and Class (or Section).
 - has prerequisite for:** Connects Course and Course.
 - registered by:** Connects Class (or Section) and Registration.
- Annotations:**
 - Arrows point to specific attributes: Student ID, Course Code, Course Number, Days / Times, and Registration Date.

this example we:

- move multi-valued Class attributes into their own entity – Class Lecture
- resolve the M:M relationship between Student and Class
- resolve the recursive Course to Course M:M relationship
- move redundant Department attributes in Course up into a new Department entity
- move Registration Status into a reference entity



World's shortest course on normalisation

Unnormalised (UNF or ONF)

- Contains multivalued attributes (a “repeating group”)

First Normal Form (1NF)

- Repeating attributes moved *down* to a dependent Characteristic or Associative entity (create a new dependent entity if necessary.) This makes data “reportable.”

Second Normal Form (2NF)

- Only applies to dependent entities
- No attribute in a child entity is really a fact about a parent (or grandparent or...)
- That is, no Characteristic or Associative entity redundantly contains facts from its parent(s) – if it does, move the fact(s) *up* (create a new parent entity if necessary)

Third Normal Form (3NF)


- If any entity redundantly contains facts from a related (non-parent) entity, move the fact(s) *out* to the other entity (create a new entity if necessary)

BCNF (Boyce-Codd NF – “3.5NF”)


- Not an issue if you keep your wits about you

Fourth and Fifth Normal Form (4NF, 5NF)

- “Large” (3-way or more) associatives need to be broken down into more granular entities



For reference – Contextual, Conceptual, & Logical models



Agree on context or “big picture”

- The scope in terms of topics or subjects that are in or out, plus core terms and definitions
- May be a simple block diagram of topics/subjects, or primarily textual (a list)
- Optional – not necessary on smaller projects

Agree on basic concepts and rules


- Ensures everyone is using the same vocabulary and concepts before diving into detail
- Overview: main entities, attributes, relationships, rules
- Lots of M:M relationships
- Relationships show cardinality
- No keys
- Few or no reference entities
- Unnormalised – most M:M relationships unresolved, many attributes will be multi-valued, redundant, and non-atomic
- Verified directly by clients plus other techniques: Use Cases...
- A “one-pager”
- 20% of the modelling effort

Full detail for physical design

- Provides all detail for initial physical database design and requirements specification
- Detailed: ~ 5 times as many entities as the conceptual model
- M:M relationships resolved
- Relationship optionality added
- Primary, foreign, alternate keys
- Lots of reference entities
- Fully normalised – no multi-valued, redundant, or non-atomic attributes. All attributes defined and “properised”
- Verified by other means: sample data, report mockups, scenarios, ...
- May be partitioned
- 80% of the modelling effort

My most plagiarised
slide ever!

Self-study exercise – from conceptual to logical



This is unnormalised – it contains multi-valued (repeating,) redundant, and constrained attributes.

First, identify the attributes that are "correct" – they are base attributes of the entity they are in. ✓
Then, normalise it to 3NF (Third Normal Form) by identifying and dealing with attributes that are:


- Multivalued, and need to be moved *down* to a dependent entity. (1NF) → ↗
- Redundant and need to be moved *up* to a parent (or higher) entity. (2NF) ↶
- Redundant or constrained and need to be moved *out* (sideways) to a related but non-parent entity, or to a reference entity. (3NF) ↗

Entity types – kernels




- The “central” objects in the model
 - “what it's all about”
 - everything else either further describes, associates, or classifies the kernel entities
- *Independent* – its existence is not dependent on another entity
 - “Does it make sense for one of these to exist on its own?”
 - is not a child of another entity
- Drawn at the top of the diagram, or subject area within a diagram

Entity types – characteristics




- Records repeating, multi-valued facts about a parent entity that have been “cast out” from the parent entity
- It “characterises” the parent entity (any type of entity)
- *Dependent* on one parent entity
- Drawn *below* the parent entity

Entity types – *associatives*




- Relates (“associates”) two or more other entities – records facts about the association (M:M relationship) between those other entities
 - sometimes so important it is discovered directly (Order, Contract, ...) and is shown on the Conceptual Model – the remainder are added on the Logical Model
 - other times it evolves from “resolving” M:M relationships
- Can associate any combination of different entity types
- *Dependent* on two or more parent entities
- Drawn *between and below* the parent entities

Associatives – notes




Associative entities often contain a date range:
Effective Date to End Date (or Last Valid Date.)
Always check for the following:

- Gaps are allowed?
- Must be contiguous?
- Overlaps are allowed?
- Effective Dates can begin in the past or future?
- Is the date range *until* or *through* the End Date? (tot en met)
- Must an End Date be specified, and if so, what format is used – “null” or “HighDate – 99991231?”
- Must the date range fit within a parent's date range?
- Do global time zones need to be handled?




Entity types – *reference or type*




- An entity that classifies or categorises other entities and/or allows the recording of standardised values for a descriptive attribute
- *Independent*
- Purpose may be served by an attribute in the Concept Model (i.e., a Customer Type attribute in the Customer entity)
- Only critical Reference entities are shown on a Concept Model (i.e., when a Reference entity ties together different parts of the model)
- Drawn beside or diagonally up from the classified entity

Graphic guidelines – the “no dead crows” principle

Draw the same kinds of things the same way every time!



Summary – entity types and conventions



"Fact" in the middle -
fine for Dimensional,
-terrible for E-R

The diagram illustrates the following tables and their attributes:


- D_PROJECTS**
 - PROJECT_KEY (PK)
 - PROJECT_DESC
 - PROJECT_OWNER
 - PCS_PROJECT_USAGE_CD
 - IN_SERVICE_FLG
- D_EMPLOYEE**
 - EMP_KEY (PK)
 - EMP_NUM
 - EMAIL_ADDR
 - EMPLOYMENT_END_DT
 - EMPLOYMENT_START_DT
 - EMP_GIVEN_NAME
 - EMP_MIDDLE_NAME
 - EMP_LAST_NAME
 - EMP_PHONE_NUM
 - FAX_NUM
 - TELEPHONE_2_NUM
 - LOGON_CD
 - GUID_TEXT
 - PICTURE_PATH_TEXT
 - SEARCH_ARGUMENT
- D_PURCHASE_MATERIAL_KEY**
 - BUYER_EMP_KEY (FK)
 - DEPARTMENT_KEY (FK)
 - PROJECT_KEY (FK)
 - PURCH_MATL_PROFILE_KEY (FK)
 - WAREHOUSE_KEY (FK)
 - METAL_SCHRGG_CD
 - PRICING_DT (FK)
 - DATE
 - TOTAL_PURCHASE_AMT
 - PURCHASE_PRICE_AMT
 - STD_COST_PRICE_AMT
 - ORDERED_QTY
 - ORDERED_AMT
 - ITEM_PURCHASE_PRICE_AMT
- D_PURCH_MATL_PROFILE_KEY**
 - PO_STATUS
 - DELIVERY_STATUS
 - ORIGINAL_CURRENCY_CD
 - TAX_TYPE_CD
 - INSTRUMENT_CD
 - FILL_TYPE_CD
 - WARRANTY_TERM_CD
 - METAL_SCHRGG_CD
 - SHIPPING_CONSTRAINT_CD
 - METAL_SCHRGG_APPLICABLE_FLG
 - REGENERATE_FLG
- D_WAREHOUSE**
 - WAREHOUSE_KEY (PK)
 - WAREHOUSE_CD
 - PACKAGING_GROUP
 - BUS_UNIT
 - NETTABLE_WHRIS_FLG
 - REJECT_WHRIS_CD
 - SPLIT_SHIPMENT_FLG
 - WAREHOUSE_DESC
 - WAREHOUSE_TYPE_DESC_EN
 - CALENDAR_CD
 - RECEIVING_LOCATION_CD
 - STAGING_LOCATION_CD
 - ENDING_LOCATION_CD
 - WAREHOUSE_CONTROLLED_FLG
- D_DEPARTMENT**
 - DEPARTMENT_KEY (PK)
 - DEPARTMENT_CD
 - CALENDAR_CD
 - CALENDAR_DESC
 - CITY_NAME
 - DEPARTMENT_DESC
 - DEPARTMENT_TYPE_EN
 - FINANCIAL_BARR_COMPANY

Relationships:

- D_PROJECTS** to **D_PURCHASE_MATERIAL_KEY**: 1:M relationship on **PROJECT_KEY**.
- D_EMPLOYEE** to **D_PURCHASE_MATERIAL_KEY**: 1:M relationship on **BUYER_EMP_KEY**.
- D_DEPARTMENT** to **D_PURCHASE_MATERIAL_KEY**: 1:M relationship on **DEPARTMENT_KEY**.
- D_WAREHOUSE** to **D_PURCHASE_MATERIAL_KEY**: 1:M relationship on **WAREHOUSE_KEY**.
- D_PURCHASE_MATERIAL_KEY** to **D_PURCH_MATL_PROFILE_KEY**: 1:M relationship on **PURCH_MATL_PROFILE_KEY**.


Let's have a quick look at Dimensional Models and their connection to ER-Models

Two sides of the house




Dimensional models

- Used to model and implement data structures for business intelligence / analytics tools. (We'll use the terms *dimensional model* and *star schema* interchangeably.)
- A central *Fact* (what you want to analyse) surrounded by *Dimensions* (how you want to select, filter, group, etc. the *Facts*)
- Each Dimension has a *primary key* which is referenced by a *foreign key* in the *Fact*, so the Fact-Dimension relationship is *M:1*




Dimensional models

- The *Fact* is measurable, quantitative data about a business event you wish to analyse.
- *Dimensions* are descriptive attributes related to the *Fact* data.
- *Dimensions* are how you want to analyse the *Fact* – how you want to “slice and dice” it.
- The *Fact* is usually an associative or characteristic entity from somewhere quite “low” in the ERD
- The *Fact* will usually include a “count” of something, even if the value is implicitly “1” - “dollars” or “hours” or “units” or ...
- The *Dimensions* are “clusters” of the *Fact*’s parents, grandparents, etc. entities
- Any combination of *Dimensions* can be used in a query, and the same *Dimension* will appear in many dimensional models




From ERD to Dimensional Model

- Any parent (or grandparent or...) entities that are encountered following M:1 relationships from the Fact are *possible* Dimensions
- Any entities that are 1:M or M:M from the Fact cannot be Dimensions without “faking” the data
- Additional Dimensions not in the original structure (e.g., Time Period) can be added
- Essentially, a basic Dimensional Model (no snowflakes) collapses an ER model to a two-level structure with a 1:M relationship between each Dimension and the Fact



A Dimensional Model is essentially a "flattened" ERD




Facts and Dimensions
are "collapsed" or "flattened"
portions of the E-R model.

From ERD to Dimensional Model


Various methodologies are available, but they can

- raise client expectations
- suggest Dimensional Models not possible with current data


We start with an ERD drawn top-down...




The Fact will be some low-level associative or characteristic



Entities at the far side of a 1:M can't be included



Entities "flattened" into Dimensions



Self-study Dimensional Modelling exercise

Jim's sister-in-law June has just returned from a BI conference, and she has Jim all wound up about building a query database so he can analyse sales (purchases by customers.)

Construct a dimensional model for Jim, using the following E-R model as a starting point. At this point, don't worry about individual attributes – just which entities would collapse into which fact or dimension. A few notes:


- Jim's has grown to a nationwide chain, with stores in many regions. Most regions cover one or more states, although some regions only cover part of a state (e.g., Northern California and Southern California). Each store is in a single city, though, and each city is in only one region.

- The layout of stores (Sections, Aisles, Store Categories, etc.) varies widely across the stores.
- The “Store Category” indicates if the store is a mall location, streetfront, “captive” (contained within another retail outlet,) etc. Web sales are not a factor.


Jim is especially interested in how the same Title sells depending on where in the Store it is displayed, because the same Title might end up in different Sections. He also wants to look at Sales by Store, Region, Artist, Publisher, Supplier, Category, ... well, just about everything! You'll have to decide what's possible and then be prepared to explain it to Jim!

Slides 90-96 are a self-study exercise – skip ahead to 97.

Dimensional Modelling exercise




Dimensional Modelling exercise




Dimensional Modelling exercise



Dimensional Modelling exercise



Solution: Dimensional Modelling exercise



This slide left blank to maintain balance in the universe

Interesting structures

➔	<i>Fundamental and Advanced Topics</i>
1. Introduction and Level-set	• Issues, Principles, Hands-on Case Study • Essentials of Concept Modelling • Transition from Conceptual to Logical, and Logical to Dimensional
2. Interesting Structures	• Types vs. Instances • Recursion, Subtyping, & Generalisation • Meeting New Requirements
3. Modelling Time, History, & Change	
4. Rules on Relationships and Associations	• Multi-way Associatives & Complex Rules • Advanced Normal Forms (4NF & 5NF)
5. Presentation Techniques for Data Modellers	• Core Techniques for Presenting • A Real-life Example

★	<i>Topics</i>
	• Types vs. Instances • Attribute vectors • Recursion and generalisation, in general • Recognising lists, trees, and networks, and modelling them with recursive relationships • Generalisation (subtyping) - when to use it, and when not to • Modelling difficult rules

Exercise: *libraries and bookstores*

Your local library and your local bookstore share some obvious similarities:

- Libraries loan books to cardholders (what the library calls a customer) and bookstores sell books to customers. Customers get to keep their purchases, but cardholders have to return whatever was loaned to them within a stated time period.

- Bookstores and libraries both keep track of all transactions (“purchase” or “loan”), but:

- the library always records the cardholder for the transaction
- the bookstore only records the customer for the transaction if they belong to their “frequent buyer” program.

Some miscellaneous points:


- Purchases and loans can both cover multiple items.
- Both of them use the term “book” somewhat loosely; they also deal with different Format Types - audiotapes, videotapes, CDs, CD-ROMs, and so on.
- They both call a “book” a “Title,” and when a Title is available on a specific Format Type (Softcover, Hardcover, CD, DVD, Audiobook, e-Book, etc.) they call that a “Release.”
- Both organisations care about the “Book’s” title and author
- Both organisations deal solely with “Books” (whatever you decide to call it) – they do not carry other types of products, or at least Books are all we care about.

What are the most important differences between the two models?

Build a simple data model for the library, and one for the bookstore. Make a guess at a few important attributes for each of the entities in your model.

The big difference - the "Types vs Instances" issue

Library – tracks *types* and *instances*



Bookstore – only tracks *types*



Each Copy (instance) of the same Title (type) has its own unique serial number – so we know which Cardholder has each.

Each instance of "Spare" does not have its own unique serial number – we can't tell one from another. All we know is how many we have. ("Quantity On Hand.")

Solution: differences and notes

Differences – library vs. bookstore

- loan vs. purchase: loan – we want it back; purchase – one-time
- one book is loaned many times in a library, but sold once in a bookstore
- library: cardholder (mandatory); bookstore: customer (optional)
- library – loan has a return date; bookstore – we hope there is NO return
- bookstore – has a price (but the library may sell books)
- *types vs. instances –*
bookstore: type (Title); library: each instance (Copy)


Diane McKellar –

"Get it wrong and live in pain forever!"


An Inventory Management system was selected that only tracked TYPES of Items.

During implementation it was discovered they also needed to track INSTANCES of certain Items – "Rotables." These were Items that could be rebuilt and "rotated" in and out of service.

The solution was a complex and expensive "shadow system" built in Excel.



Solution: libraries and bookstores



basket
analysis

Handling “vectors” of attributes

Vector:

A fixed number of repeating attributes

Divisional Sales (in 1,000,000s)				
Year	Q1	Q2	Q3	Q4
2020	1.45	1.37	1.40	1.67
2021	1.46	1.40	1.63	1.91
2022	2.11	2.32	...	...

Each row is a vector

Examples of vectors:

- 7 days of the week
- 4 quarters in a year
- 12 months in a year
- 52 weeks in a year
- 3 prices for a product
(store price, list price,
discount price)

Modelling vectors

<i>Flight number</i>	<i>Dep from</i>	<i>Time (24 hr)</i>	<i>Arr to</i>	<i>Time (24 hr)</i>	<i>Frequency</i>
SQ017	YVR	1225	SIN	2335+1	1 - - 4 - 6 -
SQ500	SIN	0715	BLR	0900	1 - - - 5 6 -
SQ502	SIN	2000	BLR	2155	1 2 3 4 5 6 7
SQ501	BLR	1015	SIN	1720	1 - - - 5 6 -
SQ503	BLR	2310	SIN	0605+1	1 2 3 4 5 6 7
SQ018	SIN	0950	YVR	1105	1 - - 4 - 6 -

Above are some of the flights you need to know about in order to travel from Vancouver (YVR) to Bangalore (BLR) via Singapore (SIN) on Singapore Airlines (SQ.)

“Frequency” indicates which days of the week the flight operates by using a string of 7 characters, with position “1” representing Monday, “2” Tuesday, through to position “7” indicating Sunday.

If the number is present, the flight operates on that day. If a dash (“-”) is present, the flight does not operate on that day.

SQ017 operates Monday, Thursday, and Saturday, and *not* on Tuesday, Wednesday, Friday or Sunday.

SQ502 operates every day of the week.

Build some alternative data models to record this subset of the flight schedule.

Make note of some of the decisions you have had to make in preparation for class discussion.

Modelling vectors – solution


Flight	
Airline Code	
Flight Number	
Departure Airport	
Departure Time	
Arrival Airport	
Arrival Time	
Operates Monday Flag	
Operates Tuesday Flag	
Operates Wednesday Flag	
....	

SQ
017
YVR
1225
SIN
2335+1

Row-wise solution

- + easy to display data
- inflexible
- queries more difficult

This is an example of the "cast out and classify" pattern




Column-wise solution

- + flexible
(change size of vector)
- + Many queries are much easier

Many "interesting" physical implementations are possible.

Recursion


- ✓ When one entity occurrence can be related to another occurrence of the same entity type
- ✓ Also know as a "self-referencing" relationship
- ✓ Three variations:
 - 1:1
 - 1:M
 - M:M
- ✓ Often involves *generalising*





With your colleagues, discuss situations in which you could use recursive 1:1 relationships – they are not common. 1:M and M:M relationships recursive relationships are very common.

Recursion - recognising the data structure



Items arranged in a linked list:



Items arranged in a tree:




Items arranged in a network:



Note - if the same two Items could be linked multiple times, we would need an Item Link SID in the primary key


Recursive 1:1 relationships



- ✓ The train above is an example of a “linked list”
- ✓ A linked list can be handled with a recursive 1:1 relationship
- ✓ All the items in the list must be generalised into the same type of entity (or a supertype with multiple subtypes - “Rail Car” would subtype into “Freight”, “Locomotive”, “Passenger”, etc.)
- ✓ The foreign key can either “point ahead” or “point back” – depends which end you add new instances from
- ✓ As always, the recursive relationship is “fully optional”
 - the first car doesn't have a car in front of it
 - the last car doesn't have a car behind it.

Recursive 1:M relationships

Partial Organization Chart - 1999/07/12





- ✓ The organisational structure is a “hierarchy”
- ✓ A hierarchy can usually be handled with a recursive 1:M relationship
- ✓ Again, this requires all the items to be generalised into the same type of entity
- ✓ The foreign key must be at the “Many” end, so the child “points to” the parent
- ✓ As with all recursive relationships, this is “fully optional”

“Org Unit” Sample Instance Table		
Org. Unit ID	Parent Org. Unit ID	Unit ID
0012	null	
0251	0012	
0064	0012	
0368	0012	
0072	0251	
0098	0251	
0391	0251	
0186	0391	
0129	0391	
0009	0064	
0267	0064	
0845	0064	
etc.	etc.	

Recursive M:M relationships

- ✓ This project plan is a “network”
- ✓ A network is handled with a recursive M:M relationship
- ✓ This resolves to an associative entity linking two different instances of the same entity






**“Task Precedence”
Sample Instance Table**

Preceding Task ID	Following Task ID
0001	0005
0001	0002
0001	0003
0005	0012
0005	0004
0002	0004
0003	0004
0012	0022
0004	0022
0022	0007
0022	0018

Future-proofing – "Avoid fixed hierarchies"


If we implement this model, what will go wrong?



This revised model is *far* more flexible.

This is an example of the "generalise, recur, and classify" pattern.


Supertypes and subtypes



- Breaks an entity down into two or more 'subtypes', or generalises two or more into a single 'supertype'
 - common relationships and attributes go into *supertype*
 - unique relationships and attributes go into *subtype*
- Subtypes are mutually exclusive and mandatory – there is exactly one subtype instance for each supertype
- a.k.a., generalisation-specification, or gen-spec

Generalisation vs. subtyping



- ✓ “Generalisation - Specialisation” is typical O-O terminology; “Supertype - Subtype” is typical E-R terminology. Gen-spec.
- ✓ Generalise whenever two or more entities, each with their own *distinct* attributes and relationships, also *share* other attributes and relationships
- ✓ Automobile, Aircraft, and Vessel have common attributes that could be generalised into Vehicle...
- ✓ ...or, Vehicle could be sub-typed into Automobile, Aircraft, and Vessel, with the same outcome
- ✓ Note that it's common for a subtyped entity to also be classified by a *type* or *reference* entity. In this example, Vehicle Type Code is the “subtype discriminator.”




Subtyping - alternative formats

Two different diagramming approaches are widely used -

- “Box in box”
(e.g., Oracle modelling tools)
- Generalisation hierarchy
(e.g., most ER- or UML-based
modelling tools)




Don't confuse “roles” and “subtypes”



- ✓ A “role” structure is used when there are two or more different “roles” that an entity type can take on.
- ✓ *Similar* to subtyping
 - unique attributes and relationships go in the role entity
- ✓ *Different* from subtyping
 - the roles are *not* mutually exclusive
 - the parent does not necessarily have to take on any role

An example with supertype, subtypes, and roles


Note that some roles are valid for one *subtype* (Person or Organisation.)
If a role is valid for both, we connect it to the *supertype* (Party.)




When to generalise

Generalisation is good if generalised items -

- ✓ are fundamentally the same
- ✓ share common facts and behaviours




Don't generalise too soon



- ✓ Confuses the client
- ✓ Reduces the chance of discovering "specifics"
- ✓ Specifics first, generalisation later

Combining recursion and generalisation

- ✓ Business rules can often be handled by a recursive relationship involving supertypes and subtypes:



Meeting a new requirement...


Confirm and extend the model:

- ✓ discover new requirements, using a variety of techniques
e.g., look for multi-valued attributes

Philosophy

- ✓ don't dive in – start simple, add detail in layers
- ✓ start out in “natural language”

See example
on the next page from our
“Good entity?” exercise




Example – meeting a new requirement

Whether facilitating or presenting,
use all three Learning Styles -
Visual, Auditory, and Kinesthetic.

In this example we used:

- Pictures (E-R Model) - *Visual*
- Words (Definition and Assertions) - *Auditory*
- Examples (Instance model) - *Kinesthetic*



If time permits, an exercise with subtyping, recursion, & rules

Extend the model so that it can record these additional facts:

1 – Organisational structure

An Organisation must be one of the recognised types, such as Corporation, Partnership, or Society.

An Organisation may be made up of multiple Organisation Units (an internal subdivision,) each of which might break down further into lower-level Organisation Units, and so on.

Each Organisation Unit has only one parent Organisation Unit. Some Organisation Units have no parent Organisation Unit, because they depend directly on an Organisation. That is, they are the highest level of Organisation Unit.

2 – Rules on Organisational structure

Each Organisation Unit is of a specific type, such as division, department, area, team, section, etc. Only certain relationships between types are valid. E.g.,


- a division can contain a department, but a department cannot contain a division.

- a team can be contained within an area or a division

Note – only certain types of Organisation Units can be immediately subsidiary to an Organisation, but we won't model that constraint at this time.

3 – Organisation ownership

An Organisation may have multiple owners, each of which could be another Organisation or a Person.



Key point:


State the constraint or fact you are trying to model in plain language before drawing the model.

Exercise work area

Solution

Assertions:

1. Each Party is either a Person or an Organisation
2. Each Organisation must be classified as one Organisation Type and Each Organisation Type may classify one or more Organisations
3. Each Organisation may divide into one or more internal Organisation Units and
4. Each Organisation Unit must be part of exactly one Organisation and Each Organisation Unit must be classified as one Organisation Unit Type and
5. Each Organisation Unit Type may classify one or more Organisation Units
6. Each Organisation Unit may control one or more other Organisation Units and
7. Each Organisation Unit may be controlled by one other Organisation Unit
8. The controlled and the controlling Organisation Units must be part of the same Organisation
9. Each Organisation Unit Type may control one or more other OU Types and Each Organisation Unit Type may be controlled by one or more other OU Types
10. The Organisation Unit Type of the controlled Organisation Unit and the Organisation Unit Type of the controlling Organisation Unit must be a pair found in Organisation Unit Hierarchy Rule
- Each Organisation may be owned by one or more Persons or Organisations (which is to say one or more Parties) and Each Party may own one or more Organisations
10. A Party may own an Organisation multiple times over a period of time



Exercise: stock exchange trading

Please build a data model from the following facts:

Companies issue shares in various stocks. For instance, Algonquin Industries has issued common stock, preferred A stock, and preferred H stock.

Each stock may be listed on multiple stock exchanges.

For instance, Algonquin's common stock is listed on the Vancouver and Toronto exchanges, but its preferred stocks are only listed in Vancouver.


When a customer wishes to buy or sell shares of a particular stock, they place an order on one of the exchanges. The order says, in effect, that "I am offering to buy (or sell) X quantity of stock Y for price Z for the next W days". If it is a sell order, the customer must also (by law) indicate if they are short selling (Short Sale Flag is set)

The Automated Trading System matches up buy and sell orders if the prices are within certain parameters. A complex algorithm determines the actual price of the sale. Note that an order may not be satisfied all at once (i.e., with one sale). For instance, an order to sell 10000 shares may be matched with a buy order for 5000 shares, another for 3000 at a later time, and it may expire before the remaining 2000 shares sell.

Build an initial E-R data model illustrating all the relevant entities, and their relationships.

1. Identify the main entities
2. Agree simple definitions
3. State assertions that describe the scenario
4. Arrange entities by dependency
5. Add and name relationships
 - name the relationship first
 - only then add cardinality (ones and manys)


Sample Instance Model




Assertions and clarifications

- Each Company may issue one or more Stocks and each Stock must be issued by one Company
- Each Stock must be classified as one Stock Type and each Stock Type may classify one or more Stocks
- Each Stock may be listed on one or more Exchanges and each Exchange may list one or more Stocks
(The Company chooses which Exchanges to list on)
- Each Customer may place one or more Buy or Sell Trade Orders & Each Buy or Sell Trade Order must be placed by one Customer
- Each Trade Order is either a Buy Trade Order or a Sell Trade Order
- Each Buy Trade Order may be filled by one or more Sell Trade Orders and each Sell Trade Order may be filled by one or more Buy Trade Orders
- The Buy and Sell Trade Orders for a Trade must be placed by different Customers
- Each Trade must be a match of one Buy Trade Order and one Sell Trade Order
- A Listing is an authorisation to buy or sell a specific Stock on a specific Exchange
- A Buy Trade Order is an offer to buy ...
A Sell Trade Order is an offer to sell...
 - a stated quantity
 - of a specific Stock
 - on a specific Exchange
 - at a specified price
 - during a specified time period
- The matching of a Buy Trade Order and a Sell Trade Order is referred to as:
 - a Sale
 - a Match
 - a Fill
 - **a Trade**
 - a Buy/Sell Transaction...

Possible entities for the Stock Exchange model




Initial Concept Model for Stock Exchange Trading



Instead of the "hash mark" (One) and "crowsfoot" (Many) this uses the symbol Miro offered at the time – an arrowhead. Miro now supports "crowsfoot" notation.

Second iteration Concept Model for Stock Exchange Trading




Key point!

Resolving the M:M between *Stock* and *Exchange* (creating *Listing*) and

generalising *Sell Order* and *Buy Order* into

Trade Order makes the model *easier* to understand!

An important constraint to check for




When there are *two* paths up to the *same parent* entity always check if the two paths must lead up to:

- the *same* parent entity
- *different* parent entities
- either the *same* or *different* parent entities (*it doesn't matter*)


We must check the paths from

- Trade to Customer - *different* Customers
- Trade to Listing - the *same* Listing

Don't generalise too soon! Specifics first, generalisation later



Complete solution: stock exchange trading




Constraints

- Buy Trade Order and Sell Trade Order involved in a Trade must reference the same Listing. Using fully inherited keys for Trade Order and Trade would make this easier to enforce.

Modelling time, history, and change

➡	Fundamental and Advanced Topics	★	Topics
1.	Introduction and Level-set • Issues, Principles, Hands-on Case Study • Essentials of Concept Modelling • Transition from Conceptual to Logical, and Logical to Dimensional		
2.	Interesting Structures • Types vs. Instances • Recursion, Subtyping, & Generalisation • Meeting New Requirements		
3.	Modelling Time, History, & Change		
4.	Rules on Relationships and Associations • Multi-way Associatives & Complex Rules • Advanced Normal Forms (4NF & 5NF)		
5.	Presentation Techniques for Data Modellers • Core Techniques for Presenting • A Real-life Example		
			• Issues in modelling history • History vs. audit trail • “What’s wrong with this model?” • Modelling for “as of” reporting • Physical time and Business time • Time-specific considerations

History in data models



- ✓ Ability to record previous values for attributes and relationships
- ✓ Ability to record history, with care, gives ability to record future cases

Issues

- E-R diagramming techniques don't support time-based rules, e.g., "1:1 at a point in time, 1:M over time"
- Two types of change – business data change, and data correction
- The client might *initially* say they don't need history, but...


Exercise: a flawed model including time dependent data

Jim needs to track List Price and Store Price over time for each Title the store carries. It is very important for him to be able to list the history of changes to either price, and the date, so that pricing information can be compared to sales figures.

Either or both prices may change at any time. Often, a price change is known before it is due to take effect.


The model to the right has several flaws – try to list at least five.

Then, construct a model (or a few alternatives) that will support Jim's needs.




A spot for you to think about that "Title Price" model


Solution – what was wrong with the flawed model?



Best solution: time dependent data



Future-proofing – "Avoid a fixed number of repeating attributes"




This model shows two types of Prices – List and Store.
Tomorrow a third will arise...
and a fourth and a fifth...

Data modellers only know
three numbers –
0, 1, and Many (One or More.)
We don't recognise 2.

This revised model offers greater flexibility and supports richer queries.

This is an example of the
"cast out and classify" pattern.



Two important time concepts

Logical Time	Physical Time
✓ Effective date/time, Start date/time, Begin date/time, etc. ✓ Time that data reflects the intent of the business at the time of update ✓ Reality Remember • Can be updated	✓ Recorded date/time, Transaction date/time, Update date/time, etc. ✓ Time when a record was written to the database ✓ Representation Remember • Cannot be updated

Wrong – with developments like Sarbanes-Oxley, we *don't change* stored data, we *add new* records.

Change and correction

The situation...

Employees are given a credit limit, which is checked whenever they attempt to make a purchase to ensure that the purchase amount does not exceed the credit limit.

Purchases are approved or rejected based on the credit limit; a record of the transaction is always made.

The credit limit changes over time, and changes to the credit limit can be entered into the system before or after the effective date of the change.


Mistakes in entering data about credit limits are common, leading to frequent corrections. Currently, the business uses a “correct in place” approach and erroneous values are overwritten with correct ones.

The problem is that after an erroneous credit limit data has been corrected, it's often difficult to see in the data why a particular transaction was approved or rejected. We need a data model that will resolve this by recording pre-corrected and corrected data.

For example...

- Rochelle, a new employee, was given a credit limit of \$5000, effective Sept 01 2019. This was entered into the system on Sept 04 2019, with an End Date of Aug 31, 2020.
- Rochelle's attempted \$3000 purchase on Sept 03 2019 was rejected.
- A year later, Rochelle's credit limit was raised to \$10,000 effective Sept 01, 2020 with no end date specified (that credit limit was to be in place indefinitely.) However, the new limit was mistakenly entered into the system as \$100,000 on Aug 22, 2020.
- That error was corrected on Sept 15 2020, when the credit limit was revised to \$10,000.
- Before the correction, Rochelle's \$17,000 purchase on Sept 08, 2020 was approved.
- In late September of 2020, auditors were not able to understand why the \$3000 Sept 2019 purchase was rejected, nor why the \$17,000 Sept 2020 purchase was approved.

Change and correction solution



	Initial	Erroneous	Corrected
Amount	\$5000	\$100,000	\$10,000
Effective Date	2019-09-01	2020-09-01	2020-09-01
End Date	2020-08-31	9999-12-31	9999-12-31
Updated End Date	null	null	null
Recorded Date	2019-09-04	2020-08-22	2020-09-15
Corrected Date	null	null 2020-09-15	null
Notes	Effective and End Dates record business intent; Recorded and Corrected Date reflect database activity.		
	An open-ended record has an End Date set to a "high date" value. If a specific End Date is later set, it is recorded in Final End Date if you don't want to overwrite End Date.		
	The Corrected Date of the corrected (previous) record is set to the Recorded Date of this correcting record.		

Four approaches to time dependent data

Example	Notes	Business Date Time	Database Date Time
Stock Listing - Current Price	"no past, "no future" "Memento"	X	X
Pressure Reading - Measurement Value Recorded Date Time	Instantaneous logging	X	✓
Program Enrollment Effective Date Time End Date Time First Valid Date Time Last Valid Date Time	the norm - for low-risk data	✓	X
Credit Limit Amount Effective Date Time End Date Time Recorded Date Time Corrected Date Time Final End Date	for "as-of reporting." Risky or regulated data "Temporal" DBs	✓	✓

Show only the **current state** of the entities

- 1 • No history, no future

Show how the data **was recorded** at various times in the past

- 2 • Shows database activity -
physical DB update dates
• Done "by accident" in older systems that used system clock as effective date

3

Show how the data **was intended** at various times in the past

- Shows business intent - effective/end dates
- This is the most common approach

4

Show both how the data **was intended**, and how the data **was actually recorded** at various times in the past

- Shows intention and correction
- Supports "as of" reporting

Time dependent data – key points

- Facts that change independently should be recorded independently
- Never name the entity “History” – it probably includes present and future values
- Distinguish between
 - *business* Effective Date
 - *database* Recorded Date
- It's tempting to put “Effective Date” in the key, but it might change, and might prevent two records from having the same Effective Date even when that makes sense
- Be sure to define what End Date / Expiry Date date *means* (“tot en met”)
- Capture the need (the “reality”) *first* in the model, *then* factor in performance considerations
- You might need to consider time zones
 - GMT / UMT / UTC
 - Local offset

Rules on relationships and associations

➔	<i>Fundamental and Advanced Topics</i>
1.	Introduction and Level-set • Issues, Principles, Hands-on Case Study • Essentials of Concept Modelling • Transition from Conceptual to Logical, and Logical to Dimensional
2.	Interesting Structures • Types vs. Instances • Recursion, Subtyping, & Generalisation • Meeting New Requirements
3.	Modelling Time, History, & Change
4.	Rules on Relationships and Associations • Multi-way Associatives & Complex Rules • Advanced Normal Forms (4NF & 5NF)
5.	Presentation Techniques for Data Modellers • Core Techniques for Presenting • A Real-life Example

★	<i>Topics</i>
	• V-A-K – using all the tools to build a model • Multi-way associations and complex rules • Advanced normal forms – resolving circular relationships and cyclic dependencies

Four key points about complex associations

1. You can't tell whether a model is correct or not simply by inspecting it – you must have business involvement

This gives rise to the other three points...


2. You must draw the model in a top-down fashion (or other systematic approach) so you can actually see dependencies
3. You must state your assumptions or understanding in narrative form as *assertions*, using terms (entity names, relationship names, and attribute names) from the data model
4. You must *illuminate* the data model by using sample data, schematic diagrams, scenarios, or some other understandable form

A quick exercise...

1. The company decides which Items will be carried at which Stockrooms.
2. The company qualifies Suppliers to provide specific Items. (A Supplier can be qualified to provide multiple Items, and an Item may be provided by multiple Suppliers)
3. The company enters into a Supply Contract with qualified Suppliers for each Item they will provide to a specific Stockroom.

Will this model satisfy the business constraints?

If not, identify specific problems and develop a better model




There are *insert anomalies* -
Can't record the
independent Supplier-Item
relationship without including
Stockroom –
"~~Stockroom #99999999~~" – a
~~dummy Stockroom~~


Can't record the
independent Stockroom-Item
relationship without including
Supplier –
"~~Supplier #99999999~~" – a
~~dummy Supplier~~

There are also
delete anomalies

First, record independent relationships



Then, associate the associates





Must go up to the
SAME Item, so
Item IDs must be
the same!



Or use one Item ID
as part of two
foreign keys

Fourth Normal Form and Fifth Normal Form

4NF and 5NF are violated when a 3-way or higher order associative entity should be broken down and made more granular.



More possibilities...



Single-valued attributes make data more "reportable."

Non-redundant – each attribute value is recorded only once.

Associative entities are as granular as necessary to reflect business rules.

This slide added to maintain balance in the universe

Presentation techniques for data modellers

➡	<i>Fundamental and Advanced Topics</i>	
1. Introduction and Level-set	• Issues, Principles, Hands-on Case Study • Essentials of Concept Modelling • Transition from Conceptual to Logical, and Logical to Dimensional	
2. Interesting Structures	• Types vs. Instances • Recursion, Subtyping, & Generalisation • Meeting New Requirements	
3. Modelling Time, History, & Change		
4. Rules on Relationships and Associations	• Multi-way Associatives & Complex Rules • Advanced Normal Forms (4NF & 5NF)	
5. Presentation Techniques for Data Modellers	★	Topics
• Core Techniques for Presenting • A Real-life Example	• A learning experience • The storyboard • A demonstration • Five key techniques	

“Here's the new Customer model.
How do you like it?”

"So what?" "Obvious!" "Yawn."
"ZZZZZZZZZZZZZZZ...."

“You’re dyin’ out there, kid! I want you to drag them through all of the pain and misery of our current files and databases. *Then* show the new model.”

“Let's try answering some important questions using the current model, and then the new model”


“Fantastic!”

“When do we get it?”

“Do you need funding? We can help!”



It's a story, so storyboard it



➔	<i>Building the storyboard</i>
1. Draw 5 "bubbles" 2. Fill in the last (your "closer" - the purpose) 3. Fill in the first (your "hook") 4. Fill in the middle ones (the "body") – add or subtract bubbles as needed 5. Allocate details to bubbles 6. Iterate until it flows and builds properly <i>Only include detail that matters!</i>	

Presenting data models

Try not to call it a *data model*

- I often call it a "world view"
- Or... "This is how Application XYZ sees the world."

Start simple, and add details in layers

- Begin with two or three fundamental things
- Work "across" the model, not a "deep dive" in one area
- Draw the model on a whiteboard as you speak to it
- Save detail like optionality until later, and primary/foreign keys until *much* later

Speak exclusively in the language of the *business*



- Don't use terms like entity, relationship, attribute, optionality, supertype, subtype, recursion, etc. Remember, you're describing a *business*, not a *database*.
- Point to the relevant entity while addressing a concept
- Someone overhearing your presentation should not realise you are presenting a data model

Make it real


- Back it up with sample data, queries, and scenarios
- Identify specific business issues or opportunities, and show how the data model helps

We'll now walk through a successful data model presentation, followed by a discussion of key points


Storyboard & 1st point for the "Flight Data Model" presentation





"Flight Data Model"




"Flight Data Model"



"Flight Data Model"



"Flight Data Model"



The five techniques that really matter

Technique	Why?	How?
1 Organise their minds to receive the presentation	- Otherwise, you're just "noise" "Why is this person telling me these things?"	"Here's the point I want to make." "This is why you care, and how I know." (even if it's obvious) "These are the caveats and limitations." "This is how I'll make my point." (storyboard!)
2 Big picture first	- Provides context and perspective - Makes subsequent detail understandable	- Show contextual data model first, build up detailed models later - Process context first, process flow later - Describe 5 problem areas first, specifics of each area later
3 Do it live	- Focuses, demands that they watch - Involves them / you - It means 'attending has value'	- Use memory triggers, not a script - Build up content progressively on white board, flip chart, or screen - Add brainstorming, discussion, or questions - Have them physically "do stuff"
4 Present information in various forms	- Adds interest - Different forms have different strengths	- Supplement PowerPoint slides with flip charts, white boards, Post-Its, handouts, etc. - Use props – the thing itself, not a description - Use visual, auditory, and kinesthetic approaches
5 Show, then tell	- Point is more meaningful if experienced firsthand - Saves time, simplifies	- Scenario / example first, then concept / abstraction - Problem first, solution second - Thing first, description / discussion second

Other courses for analysts by Alec Sharp

Working With Business Processes – Process Change in Agile Timeframes 2 days

Business processes matter, because business processes are how value is delivered. Understanding how to work with business processes is now a core skill for business analysts, process and application architects, functional area managers, and even corporate executives. But too often, material on the topic either floats around in generalities and familiar case studies, or descends rapidly into technical details and incomprehensible models. This workshop is different – in a practical way, it shows how to discover and scope a business process, clarify its context, model its workflow with progressive detail, assess it, and and transition to the design of a new process by determining, verifying, and documenting its essential characteristics. Everything is backed up with real-world examples, and clear, repeatable guidelines.

Business-Oriented Data Modelling – Useful Models in Agile Timeframes 2 days

Data modelling was often seen as a technical exercise, but is now known to be essential to other initiatives such as business process change, requirements specification, Agile development, and even big data, analytics, and data lake implementation. Why? – because it ensures a common understanding of the things – the entities or business objects – that processes, applications, and analytics deal with. This workshop introduces concept modelling from a non-technical perspective, provides tips and guidelines for the analyst, and explores entity-relationship modelling at contextual, conceptual, and logical levels using techniques that maximise client involvement.

Working With Business Processes Masterclass – Aligning Process Work with Strategic, Organisational, and Cultural Factors 3 days

This 3-day interactive workshop combines the core content from two highly-rated classes by Alec Sharp – “Working With Business Processes” and “Advanced Business Process Techniques.” This structure is popular because it gets both new and experienced practitioners to the same baseline on Clartiq’s unique, agile, and ultra-practical approach to Business Process Change. First, it shows how to effectively communicate Business Process concepts, discover and scope a business process, assess it and establish goals, and model it with progressive detail. Then, it shifts to advanced topics – specific, repeatable techniques for developing a process architecture, encouraging support for change, and completing a feature-based process design. The emphasis is always on ensuring business process initiatives are aligned with human, social, cultural, and political factors, and enterprise mission, strategy, goals, and objectives.

Business-Oriented Data Modelling Masterclass – Balancing Engagement, Agility, and Complexity 3 days

Our most popular workshop! This intensive 3-day workshop combines the core content from two popular offerings by Alec Sharp – “Business Oriented Data Modelling” and “Advanced Data Modelling.” First, the workshop gets both new and experienced modellers to the same baseline on terminology, conventions, and Clarteq’s unique, business-engaging approach. We ensure a common understanding of what a data model *really* is, and maximising its relevance. Then, we provide intense, hands-on practice with more advanced situations, such as the enforcement of complex business rules, handling recurring patterns, satisfying regulatory requirements to model time and history, capturing complex changes and corrections, and integrating with dimensional modelling. Always, the philosophy is that a data model is a description of a business, not of a database, and the emphasis is on engaging the business and improving communication.

Model-Driven Business Analysis Techniques – Proven Techniques for Processes, Applications, and Data 3 days

Simple, list-based techniques are fine as a starting point, but only with more rigorous techniques will a complete set of requirements emerge, and those requirements must then be synthesised into a cohesive view of the desired to-be state. This three-day workshop shows how to accomplish that with an integrated, model-driven framework comprising process workflow models, a unique form of use cases, service specifications, and business-friendly data models. This distinctive approach has succeeded on projects of all types because it is “do-able” by analysts, relevant to business subject matter experts, and useful to developers. It distills the material from Clarteq’s three, two-day workshops on process, data, and use cases & services.

*** *Note: two-day in-person workshops are delivered virtually as three half-day sessions via Zoom.*

Three-day in-person workshops are delivered virtually as five half-day sessions via Zoom.

Thank you – stay in touch!



Alec Sharp
Clariteq Systems Consulting Ltd.
West Vancouver, BC, Canada

- asharp@clariteq.com
- ig: [@alecsharp01](#)
- www.clariteq.com
- <http://amzn.to/dHun1o>



And most of all, if you have questions or comments...
don't be shy – send me a note!