<div align="center">

Informatics 1B. Tutorial 2
Data and Analysis: Relational Algebra

Author: Marques-Pita, 2006
Revised and issued by Kate Byrne

February 6, 2007

</div>

1 Introduction

This tutorial is the continuation of the work you started in week 3, in Tutorial 1. The focus of this tutorial is on mapping the ER model you built last week to a *relational schema*. What you should know for this tutorial is covered in section 1.3 of the lecture notes.

Question 1 — *Vessel*

Using an SQL construction, how should we define the *Vessel* relation? For attribute domains assume that SQL has the following built-in types to choose from: `integer`, `real`, `char(n)`, `boolean`, `date` and `timestamp`.

Question 2 — *Journey*

Note that the given ER diagram for Tutorial 1, in Figure 1, defines the entity *Journey* as a weak entity set.
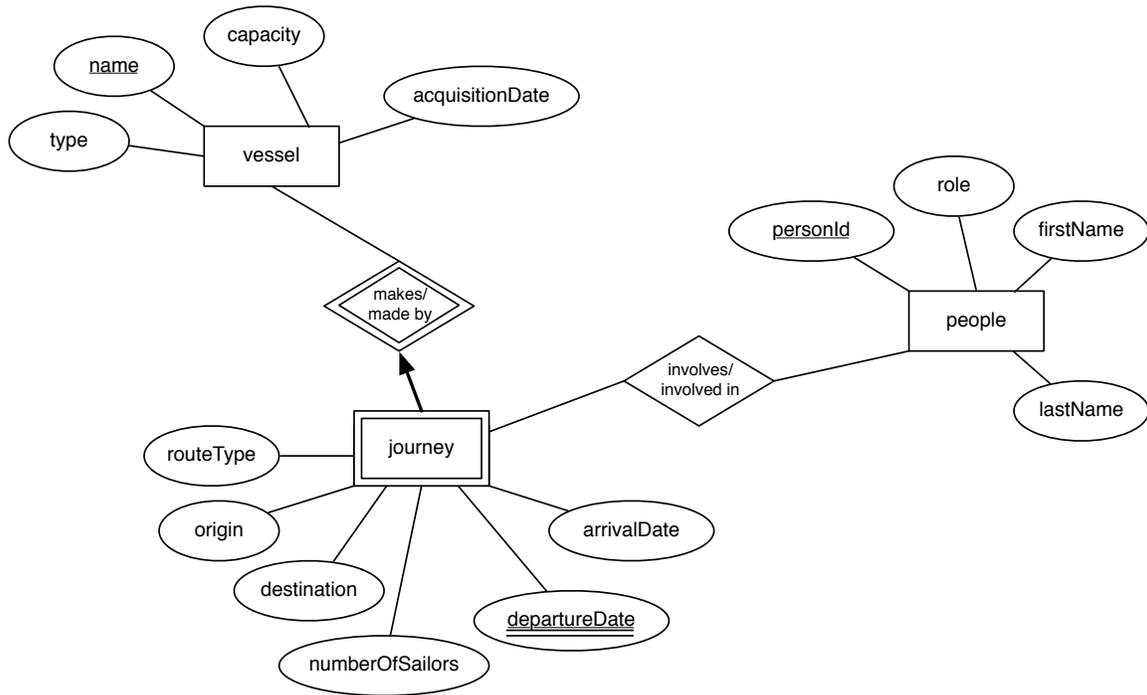
Question 2.1

Why is *Journey* a weak entity?

Figure 1: a complete model for the Guipuzcoana example

Question 2.3

What is the *identifying owner* in this weak entity set? Why?

Question 2.4

What should happen to the the stored instances of journeys linked to one vessel instance if that vessel is deleted from the *Vessel* relation?

Question 2.5

Considering the answers to the previous three questions, how do we define the relation *Journey* in SQL?

Question 3 — *People*

Question 3.1

Define the *People* relation using an SQL construction.

Question 3.2

How do we make sure the relationship between *People* and *Journey* is accounted for in the relational schema?

Question 3.3

What would happen to the design if (work on each of the following separately):

1. we want to include a field in the *People* relation to store the date an employee started working for the company;

2. employees can have different roles in different journeys — i.e. person X was the chef on journey Y but a sailor in journey Z.