


Rutgers University, Business School/Undergraduate New Brunswick
Operations Management (33:623:370:03)
 Fall 2008; Instructor: Jonathan Eckstein

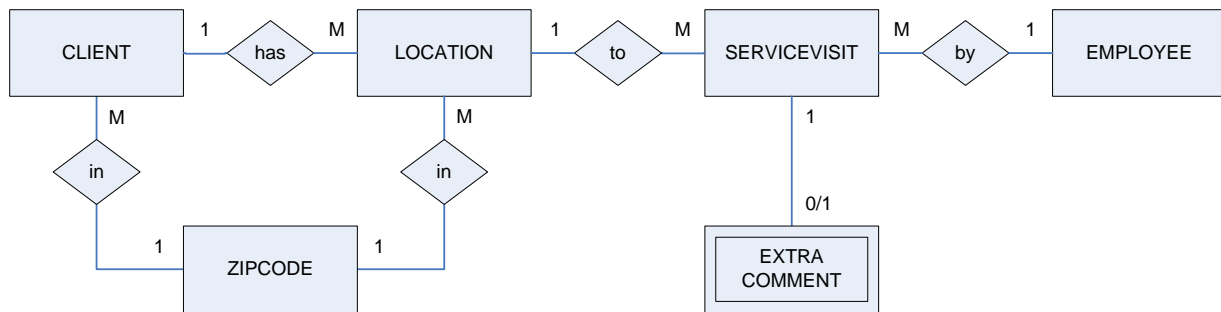
Solutions to Midterm Exam — Friday, October 17, 2008

	Q1	Q2	Multiple Choice	Total
Points Allocated	50	25	35	100
Highest Score	50	26	24	96
Mean	44.6	20.3	18.7	83.7
Median	46.5	22.0	19.5	85.0
Lowest Score	31	11	10	63
Standard Deviation	5.8	4.2	3.4	9.4
Mean as %	89.3%	81.3%	74.8%	83.7%



The average score on this exam was just below 84, which I consider good, and about the difficulty I like to see (if scores are much higher than this, I do not get enough information to base my grades on). I do not “curve” individual exams, but I do “curve” (in my own manner) your overall performance accumulated at the end of the course.

Part I, Q1: Nerd Patrol (50 points)



EMPLOYEE(EmployeeID, FirstName, LastName, HireDate)

CLIENT(ClientID, Name, ContactName, BillPhone, BillAddress, BillZip, HourlyRate)
 BillZip foreign key to ZIPCODE

LOCATION(LocationID, Name, Phone, Address, Zip, ClientID)
 Zip foreign key to ZIPCODE
 ClientID foreign key to CLIENT

SERVICEVISIT(VisitID, EmpID, LocationID, DateTimeStart, DateTimeEnd, Description)
 EmployeeID foreign key to EMPLOYEE
 LocationID foreign key to LOCATION

EXTRACOMMENT (VisitID, ExtraComments)
 VisitID foreign key to SERVICEVISIT

ZIPCODE(Zip, City, State)

This problem was fairly similar to one of the practice problems, but less complicated. Unlike that practice problem, however, it had several different tables relating to a master zip code table, more like one of the examples we studied in class. As a whole, the class did well on this problem, scoring an average of 44.6/50 points, or about 89%. This performance shows that most students have mastered the basics of database design – the most important topic in the course.

Note that when I say that a master zip code table is available, that means that you should include such a table in your design, with the proper relationships. If I say a master zip-code table is *not* available, then you can leave the table out of your design and pretend that treat *city*, *state*, and *zip* are independent fields. I dedicated class time specifically to this subject, discussing it very explicitly; I also gave examples of both kinds of designs on the homework.

I have shown synthetic keys for all the tables, since attractive alternatives were not available. Note that whenever you show a subtype in your ER diagram, the primary keys of the two tables involved must be identical. It makes sense to have a subtype for EXTRACOMMENTS, since the SERVICEVISIT table is likely to be large, and the *ExtraComments* field is both large and rarely used.

My grading scheme was as follows:

- 4 points for each of the 5 tables except EXTRACOMMENT, broken down as follows:
 - 1 point for having the table your diagram
 - 1 point for having the table in your outline
 - 1 point for having the correct primary key
 - 1 point for the remaining fields.
- 5 points for each of the 5 one-to-many relationships, scored as follows:
 - If an entire relationship was missing, misplaced, or reversed, both in the diagram and outline, you lost the full 5 points
 - If a relationship was incorrect in your diagram, but correct in your outline, you lost only 2 point
 - If a relationship was correct in your diagram, but incorrect in your outline, you lost 3 points.
- 5 points combined for the EXTRACOMMENT table and its relationship to SERVICEVISIT.
 - If you showed the relationship as a subtype in your diagram, but implemented it as a one-to-many relationship in your outline, you lost 3 points.
 - If you implemented the relationship between SERVICEVISIT and EXTRACOMMENT as a one-to-many, and depicted it consistently in your diagram, you lost just 1 point.

Here are some other common problems and point deductions:

- Treating the problem as if I had said a zip code table was not available, with *city*, *state*, and *zip* fields in both CUSTOMER and LOCATION. In this case, I deducted 7 points, or half the total for ZIPCODE table and its two relationships ($4 + 5 + 5 = 14$ points). Completely omitting zip codes and addresses from your answer (a rare problem) cost the full 14 points. If you had a correct ZIPCODE table, you lost 2 points for each other table that also had *city* and *state* fields.
- An unnecessary extra relationship between EMPLOYEE and LOCATION or EMPLOYEE and CUSTOMER. These relationships can be inferred by tracing foreign keys through the SITEVISIT table. I deducted 5 points for such errors (the same as for a single misplaced relationship).

Part I, Q2: Memory Storage Calculations

People had a bit more trouble with this section of the exam than usual. Still, the mean score was a reasonably good 20.3/25 points, or about 81%.

(a) [4 points] For widescreen,

$$(1440 \times 900 \text{ pixels}) \left(24 \frac{\text{bit}}{\text{pixel}} \right) \div \left(8 \frac{\text{bits}}{\text{byte}} \right) = 3.88 \times 10^6 \text{ bytes}$$
$$(3.88 \times 10^6 \text{ bytes}) \div \left(1024^2 \frac{\text{bytes}}{\text{MB}} \right) \approx \boxed{3.71 \text{ MB}},$$

whereas for a traditional screen,

$$(1280 \times 1024 \text{ pixels}) \left(24 \frac{\text{bit}}{\text{pixel}} \right) \div \left(8 \frac{\text{bits}}{\text{byte}} \right) \approx 3.93 \times 10^6 \text{ bytes}$$
$$(3.93 \times 10^6 \text{ bytes}) \div \left(1024^2 \frac{\text{bytes}}{\text{MB}} \right) \approx \boxed{3.75 \text{ MB}},$$

so the traditional screen image requires slightly more storage. Note that the 19-inch measurement of the screen is not part of the calculation (a few people multiplied their answers by 19, which would formally give you an answer in inch-megabytes). Aside: because of the way screens are measured along the diagonal, “widescreen” monitors also have less viewing area than traditional monitors described as having the same “size” in inches.

(b) [4 points] Note that 32 bits/pixel = 4 bytes/pixel. Then,

$$(8.5 \text{ inches}) \left(1200 \frac{\text{divisions}}{\text{inch}} \right) (11 \text{ inches}) \left(1200 \frac{\text{divisions}}{\text{inch}} \right) \left(48 \frac{\text{bits}}{\text{pixel}} \right) \div \left(8 \frac{\text{bits}}{\text{byte}} \right) \approx 8.08 \times 10^8 \text{ bytes}$$
$$(8.08 \times 10^8 \text{ bytes}) \div \left(1024^2 \frac{\text{bytes}}{\text{MB}} \right) \approx \boxed{770 \text{ MB}}.$$

A common error here was to omit one of the factors of 1200. Remember, there are 1200 divisions per inch, both vertically and horizontally.

(c) [1 point] If your answer here was incorrect, but matched your answer to (a), I did not deduct more points.

$$\left(\frac{770 \text{ MB original}}{16 \text{ MB compressed}} \right) \approx \boxed{48.2}$$

- (d) [5 points] Remember, communication line speeds are given in *decimal* bits per second, and file sizes are usually in *binary*-style bytes. You have to compensate, as explained in the course notes. For the DVD movie, we calculate

$$(4 \text{ GB}) \left(1024^3 \frac{\text{bytes}}{\text{GB}} \right) \left(8 \frac{\text{bits}}{\text{byte}} \right) \approx 3.44 \times 10^{10} \text{ bits}$$

$$(3.44 \times 10^{10} \text{ bits}) \div \left(18,000,000 \frac{\text{bits}}{\text{second}} \right) \approx 1909 \text{ seconds} \div \left(60 \frac{\text{seconds}}{\text{minute}} \right) \approx \boxed{31.8 \text{ minutes}}.$$

For the HD version, we instead get

$$(20 \text{ GB}) \left(1024^3 \frac{\text{bytes}}{\text{GB}} \right) \left(8 \frac{\text{bits}}{\text{byte}} \right) \approx 1.72 \times 10^{11} \text{ bits}$$

$$(1.72 \times 10^{11} \text{ bits}) \div \left(18,000,000 \frac{\text{bits}}{\text{second}} \right) \approx 9544 \text{ seconds} \div \left(60 \frac{\text{seconds}}{\text{minute}} \right) \approx \boxed{159 \text{ minutes}}.$$

Thus, the HD version takes over two and half hours, which is probably too long. (Also, actual times could be longer due to network header information and congestion).

Downloading the HD movie in 10 minutes would require a communication line operating at

$$\left(\frac{171,798,691,840 \text{ bits}}{(10 \text{ minutes}) \left(60 \frac{\text{seconds}}{\text{minute}} \right)} \right) \approx 2.86 \times 10^8 \frac{\text{bits}}{\text{second}} = \boxed{286 \text{ Mb/s}},$$

and possibly faster if we take into account header overhead and network congestion.

People had a lot of trouble with this last calculation, or didn't put their answer in Mb/s units, so I treated it as a 1-point extra credit exercise.

- (e) [5 points] We first note that 24 bits/sample = 3 bytes/sample. Then, without compression, we calculate

$$\left(100,000 \frac{\text{samples}}{\text{second} \cdot \text{channel}} \right) \left(3 \frac{\text{bytes}}{\text{sample}} \right) (8 \text{ channels}) \left(60 \frac{\text{seconds}}{\text{minute}} \right) \left(60 \frac{\text{minutes}}{\text{hour}} \right) \approx 8.64 \times 10^9 \frac{\text{bytes}}{\text{hour}}.$$

To get the bytes per hour with compression, we divide by 15:

$$\left(8.64 \times 10^9 \frac{\text{bytes original}}{\text{hour}} \right) \div \left(15 \frac{\text{bytes original}}{\text{bytes compressed}} \right) = 5.76 \times 10^8 \frac{\text{bytes compressed}}{\text{hour}}.$$

Finally, we calculate

$$\left(4.7 \frac{\text{GB}}{\text{disk}} \right) \left(1024^3 \frac{\text{bytes}}{\text{GB}} \right) \div \left(5.76 \times 10^8 \frac{\text{bytes compressed}}{\text{hour}} \right) \approx \boxed{8.76 \frac{\text{hours}}{\text{disk}}}.$$

There was a typo in the exam (corrected on the board), where I wrote “MB” when I meant “GB”. In that case, you would get about 0.0086 hours (for which I gave full credit).

(f) [4 points]

Attribute	Access Datatype	Size in Bytes
Zip	Text, 5 characters	5
City	Text, 35 characters	35
State	Text, two characters	2
Population	Long integer	4
Median income	Currency	8
Total		54

We multiply the 54 bytes per row by 43,000 rows, obtaining

$$\left(54 \frac{\text{bytes}}{\text{row}}\right)(43,000 \text{ rows}) = 2,322,000 \text{ bytes} \div \left(1024^2 \frac{\text{bytes}}{\text{MB}}\right) \approx \boxed{2.21 \text{ MB}}.$$

(g) [2 points] To download this file at 1 Mb/s, we just calculate

$$(2,322,000 \text{ bytes}) \left(8 \frac{\text{bits}}{\text{byte}}\right) \div \left(1 \times 10^6 \frac{\text{bits}}{\text{second}}\right) \approx \boxed{18.6 \text{ seconds}}.$$

Note that the actual time could be longer due to header information added by the various network layers, and network congestion.

Part II: Multiple Choice

This section was worth 25 points, one for each question. As usual, this was the hardest part of the test, with a mean of about 75%. Unless otherwise specified, the pages indicated below refer to the lecture notes handouts.

1. In the terminology of the study of information systems, the relationship between the terms *data*, *information*, and *knowledge* is:
 - a. They mean the same thing
 - b. Data and information become knowledge once they are transmitted through a network
 - c. Data consists of information structured to be useful in making some decision or performing some task; knowledge and information are synonymous
 - d. Information consists of data structured to be useful in making some decision or performing some task; knowledge conveys some kind of human “understanding” of information – *page 1*
 - e. Knowledge becomes data once it is stored in a computer system; information consists of data structured to be useful in making some decision or performing some task.

2. Generally speaking, information systems consist of hardware, software, procedures, databases and _____ .
- Modules
 - Queries
 - Algorithms
 - Programs
 - People – *top of page 2*
3. When did it become common for human organizations to start using information systems?
- Once electronic computers were invented around 1950
 - Once electronic computers became commonly used to manipulate business information around 1960
 - In the early 1990's: you cannot have a “true” information system without the internet, and businesses did not start using the internet until the early 1990's
 - Around 1985; before that, businesses only had data systems, not information systems
 - When human organizations first came into being in prehistory; however, the hardware and procedures used in those times were very different from today. – *see page 2, where it says “information systems existed before computers and networks”*
4. Which of the following is *not* an example of a functional area information system?
- A system used by the accounting department to maintain the accounts payable and accounts receivable ledgers
 - A system used by human resources to track employee healthcare expenses
 - A contact management system used by marketing to track customer behavior
 - An ERP system used by marketing, accounting, and operations – *page 3*
 - A transaction system used by the operations to track the movement of parts shipments.
5. Which of the following is considered a *transaction* to be tracked by a TPS?
- A customer orders a product
 - A customer returns a defective product
 - UPS picks up a shipment from your loading dock
 - All of the above – *page 3*
 - Both (a) and (b), but not (c)
6. An *intranet* typically denotes _____ .
- The portion of an organization's network that is accessible to those outside the organization
 - The portion of an organization's network that is *not* accessible to those outside the organization – *page 5*
 - A network on which most traffic is encrypted
 - A network that uses a single network layer, but more than one transport layer
 - A network that allows only one transport layer.

7. A benefit of the Sarbanes-Oxley act is that _____ .
- It requires employees to become more proficient with spreadsheets
 - It requires firms to upgrade to the most recent available software
 - It has prompted the replacement of semi-manual, spreadsheet-based information handling with more appropriate, secure information systems – *pages 6-7*
 - It requires all “legacy” systems to be replaced with ERP technology
 - It requires information systems to use foreign keys instead of repeating groups.
8. In Microsoft Access, suppose you designate a field to be a “number” with field size “double”. This field will consume _____ byte(s) of storage per record.
- 1
 - 2
 - 4
 - 8 – *memory storage calculations handout*
 - 16
9. The basic capabilities of Microsoft Access do not include _____ .
- Transaction processing
 - Analytical processing
 - Deleting records
 - Decision support – *page 11*
 - Joins.
10. If a communication link functions at 56 Kb/s, that means it transmits _____ .
- 56,000 bits per second – *pages 13 and 18-19 note that communication link speeds use decimal prefixes*
 - 57,344 bits per second
 - 448,000 bits per second
 - 458,752 bits per second
 - 56,000,000 bits per second.
11. A customer calls your customer service center, orders two products, and informs you of a change of address. Entering the corresponding data into your information system constitutes an example of _____ .
- Transaction processing – *pages 3, 11*
 - Analytical processing
 - Data mining
 - Decision support
 - All of the above.

12. The _____, originally built in the 1970's, evolved into today's internet.
- a. TELNET
 - b. ARPANET – *page 14*
 - c. Modem
 - d. FTPNET
 - e. Name server subnet.
13. In TCP, _____ .
- a. Messages may have multiple ports, and each port may have multiple sessions
 - b. Ports may have multiple messages, and each message may have multiple sessions
 - c. Ports may have multiple sessions, and each session may have multiple messages – *pages 15-16*
 - d. Sessions may have multiple messages, and each session may have multiple ports
 - e. Messages may have multiple sessions, and each session may have multiple ports.
14. *IP addresses* and *hostnames* are defined by _____ .
- a. The physical network layer
 - b. The IP network layer – *page 15*
 - c. The TCP transport layer
 - d. DHCP
 - e. NAT.
15. “128.43.12.18” is an example of a(n) _____ .
- a. IP address – *page 15*
 - b. Hostname
 - c. URL locator
 - d. Name server
 - e. DHCP request.
16. “rbsmail.rutgers.edu” is an example of a(n) _____ .
- a. IP address
 - b. Hostname – *page 15*
 - c. URL locator
 - d. Name server
 - e. DHCP request.
17. The transition from IPv4 to IPv6 will eventually be necessary because _____ .
- a. We should reduce packet overhead to improve the efficiency of the internet
 - b. With IPv4, the internet is running out of distinct IP addresses – *page 15*
 - c. With IPv4, the internet is running out of distinct hostnames
 - d. NAT and DHCP are unreliable
 - e. Use of SSH is becoming widespread.

18. "Packet switching" refers to _____ .
- a. Collecting multiple messages into larger "packets" that must take the same route through the network
 - b. Dividing individual messages into smaller "packets" that may take different routes through the network – *page 16*
 - c. Intercepting a message sent to one computer and "switching" it to another computer
 - d. Using a router to allow several computers (for example, in your home) to share the same connection to the internet
 - e. Using two different protocols to transmit the same information.
19. Which of the following is used to transmit e-mail?
- a. DHCP
 - b. FTP
 - c. NFS
 - d. SMB
 - e. SMTP – *page 16*
20. SSH is an example of a(n) _____ layer/protocol.
- a. Physical
 - b. Network
 - c. Transport
 - d. Application – *page 16*
 - e. Obsolete.
21. Which of the following is not considered a "last mile" technology?
- a. Telephone line and modem connections
 - b. DSL
 - c. Cable Modems
 - d. WiFi – *page 18*
 - e. WiMax.
22. The term "ISP" denotes _____ .
- a. An alternative to IP
 - b. An alternative to TCP
 - c. A kind of secure application protocol
 - d. A firm that provides internet connections and/or maintains portions of the internet – *top of page 18*
 - e. A secure interconnection path between two firms' networks, whose traffic cannot be detected by internet users outside those two firms.

23. One reason to keep information in a data warehouse is _____ .
- a. To keep it secure from unauthorized access
 - b. So that analytical processing queries can use the most recent data in the organization
 - c. So that you can have a backup copy in case of system malfunctions
 - d. So that analytical processing queries do not slow down your transaction systems –
pages 11-12
 - e. To share it with business partners such as suppliers and regular customers.
24. You can share a single IP address between several computers by using _____ .
- a. DHCP
 - b. IMAP
 - c. NAT – *page 15*
 - d. NFS
 - e. SSH
25. The personnel that maintain and develop a firm's information systems are part of _____ .
- a. The firm's information architecture but not part of its information infrastructure
 - b. The firm's information infrastructure but not its information architecture – *page 2*
 - c. Both the firm's information architecture and its information infrastructure
 - d. Neither the firm's information architecture nor its information infrastructure
 - e. The firm's "software".