Computer Science Graduation Exam 2018
Practice Questions - Computing Systems and Databases

Note:

The written test of the graduation exam (July or September 2018) will consist of 60 questions which will be similar (in structure and difficulty level) with those included in this booklet. For each of the three categories (Discrete Structures and Algorithms, Programming Languages and Software Engineering, Computing Systems and Databases) there will be 20 questions.

If there are unclear aspects concerning the questions and/or answers please contact the teacher(s) who proposed the questions for each section:

Computer Architecture:

- Adrian Crăciun (adrian.craciun@e-uvt.ro)

Operating Systems:

- Ciprian Pungilă (ciprian.pungila@e-uvt.ro)
- Florin Fortiş (florin.fortis@e-uvt.ro)

Databases:

- Monica Sancira (monica.tirea@e-uvt.ro)
- Daniel Pop (daniel.pop@e-uvt.ro)

Computer Networks:

- Stelian Mihalaş (stelian.mihalas@e-uvt.ro)

1 Computer Architecture

1. By translation,
   (a) a program P that written for a high-level machine M is transformed by a translator into a program P_1 that can then be executed on a lower level of that machine, without further need for P
   (b) each instruction of a program P written for a high-level machine M is transformed into a sequence of instructions for a lower level of that machine, then immediately executed, before the same process is applied to the next instruction from P
   (c) a program P that is written for a high-level virtual machine VM is transformed into a program P_1 which can be executed on a physical machine M
2. A virtual machine
   (a) is defined by the language whose instructions can be executed on the respective machine
   (b) is a model of a machine that cannot be implemented in hardware
   (c) represents a level in the multilevel organization of computer systems, used to control the execution of microinstructions
3. The von Neumann architecture is a term used to denote:
   (a) the design of a multiprocessor computer with distributed memory, that influenced the development of supercomputers
   (b) the design a computer system consisting of an arithmetic logical unit, with registers, a memory that contains programs and data, a control unit and input-output devices, whose influence can be traced in today's popular architectures
   (c) the design of a system of buses and control mechanisms that allow direct interaction between memory and input-output devices
4. Which of the following statements are consequences of Moore's law:
   (a) computing power is enjoying exponential growth
   (b) over a short period of time (18-24 months) the size of a digital component (chip) can shrink to roughly half
   (c) processor designs are considered obsolete after a short period of time (18-24 months)
5. The data path is defined as:
   (a) the connection between input-output devices and the processor
   (b) the representation of data types in the memory of a computing system
   (c) the connection trough buses between registers arithmetic-logical unit and back to registers, used for instruction execution

(d) networks interconnecting computing systems (like the Internet)

6. The program counter:

   (a) is a global variable in a program that is used to count the number of instructions executed, and measure program complexity

   (b) is a register that holds an address to the next instruction that has to be executed

   (c) is a register that holds the instruction that is being executed in a processor

   (d) is an operating system variable that counts the number of processes in execution at a given time

7. Microprogramming is

   (a) the activity of writing programs that run on microprocessors

   (b) the use of a set of instructions (microinstructions), stored in a memory, that are used to control the data path, for the purpose of executing instructions in one or several data path cycles

   (c) the interaction with input/output devices by issuing commands to controllers (e.g commands for the hard drive controller to read a word from an address)

   (d) automated generation of small programs for embedded devices.

8. Which of the following statements about microprogramming is true

   (a) new instructions can be added to the microprogram, allowing the extension of the instruction set associated with the respective architecture

   (b) software patches can correct potential design problems/bugs, etc

   (c) the instruction sets can become too complex, hard to decode, slow to execute

   (d) makes it easy to write compilers for the respective target architecture

9. A RISC architecture and its implementation is characterized by:

   (a) simple instructions, most of which can be executed in 1 clock cycle

   (b) a big number of registers

   (c) powerful instruction decoding hardware

10. Pipelines

    (a) are the wires that connect the different components of the computer

    (b) are hardware mechanisms through which multiple input-output devices (such as hard disks) can be connected to a computer and get access to the bus in the same time

    (c) separate the stages of the fetch-decode-execute cycle, thus allowing multiple instructions to be executed in the same time

11. A superscalar CPU is

    (a) a CPU design where the pipeline has multiple functional units in parallel (for the execution stage of the pipeline)

    (b) a CPU specially designed for low-power embedded applications

    (c) a system-on-a-chip design

12. A kilobyte is

    (a) 1000 bits

    (b) 1024 bytes

    (c) 2^{13} bits

13. RAID is a class of methods of organizing secondary memory that

    (a) stores data in big endian mode for easier access

    (b) ensures faster access to data on disks by using several small hard drives that are logically seen as one big hard drive

    (c) ensures protection of data by using various redundancy schemes (such as disk duplication, or error correcting codes)

    (d) gives access to a bigger address space compared to non-RAID schemes.

14. A typical use for a demultiplexer circuit is the implementation of

    (a) a 1 bit full adder circuit

    (b) a serial-to-parallel convertor

    (c) a parallel-to-serial convertor

    (d) an integer multiplication circuit

15. A programmable logic array is programmed by

    (a) using assembly programming languages

    (b) using microinstructions

    (c) blowing up fuses in the circuit in order to define formulae in disjunctive normal corresponding to boolean functions

    (d) by interpretation of special high-level scripting languages

16. The following inputs

    $I_1 = 0$, $I_2 = 1$, $I_3 = 1$,
    $A_1 = 1$, $A_2 = 0$,
    $CS = 0, RD = 0, OE = 1$

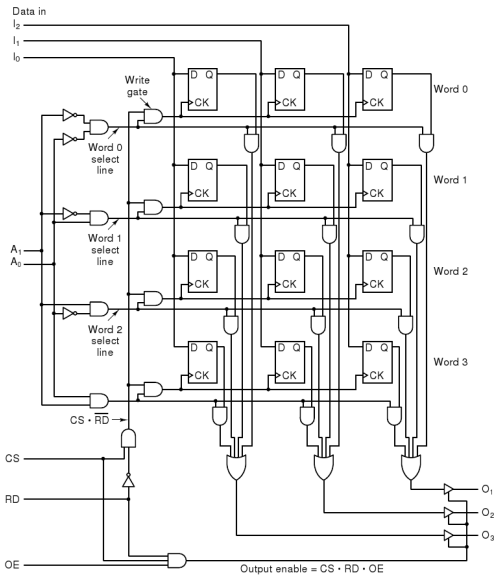

Figure 1: A 4x3 bit memory circuit.

determine the following behavior for the the circuit in Figure 1

(a) 111 is written on word 2,

(b) the content from word 2 is put on the lines O_1, O_2, O_3

(c) nothing happens, the chip is not selected

17. The Integer Java Virtual Machine and its implementation, Tanenbaum's Mic-1 can be described as

(a) a stack machine

(b) a RISC machine

(c) a CISC machine

18. For cache memories, the locality principle refers to:

(a) the fact that memory locations close to the previously used one are likely to be used in the near future

(b) the fact that memory locations that were accessed recently will likely be accessed again

(c) the position of cache levels with respect to the processor (on processor, or very near the processor)

19. Which of the following statements are true:

(a) the use direct mapped caches vs set associative caches will lead on average to fewer cache misses

(b) the use set associative caches vs direct mapped caches will lead on average to fewer cache misses

(c) the average number of cache misses does not depend on the caching mechanism used

20. Which type of branch instruction causes most delay in a typical five stage pipeline:

(a) unconditional branches

(b) conditional branches

(c) none of the branch prediction causes any delay in the pipeline

(d) both types of branch prediction cause the same amount of delay, on average

21. Register windows are used to

(a) implement efficient procedure calls on RISC machines

(b) enable or disable execution units in superscalar computers

(c) record information about the placement of windows on desktop graphical interfaces

22. Which of the following are necessary to describe the Instruction Set Architecture level:

(a) the memory model

(b) the configuration of the cache

(c) the configuration of registers

(d) data types

(e) instructions

23. What addressing mode is the most used for instructions in a RISC architecture

(a) direct addressing

(b) register addressing

(c) register indirect addressing

(d) indexed addressing

24. Which of the following situations will generate an interrupt?

(a) the hard drive has located data and is ready to transfer

(b) an overflow has occurred in the running of a program

(c) a key is pressed on the keyboard

(d) a program has run out of memory

25. Segmentation is a virtual memory technique that

(a) provides multiple address spaces
(b) is done automatically, without the need for explicit intervention from the programmer
(c) allows separation of instructions and data in the memory

26. Preemptive scheduling refers to

(a) a technique that assigns to several running processes a time slice in which they have control of the processor, to simulate parallel execution
(b) the execution of instructions before they are needed to avoid the situation where the execution units of a CPU are not used for several cycles
(c) the reservation of continuous space on hard drives such that large files will not be fragmented when stored

27. The result of the following addition

$$\begin{array}{r} 11110111 \\ +11110111 \end{array}$$

in 2's complement representation on 8 bits is:

(a) 11101110 and the result is correct
(b) 11101110 but there is an overflow
(c) 11101111 and the result is correct
(d) 11101111 but there is an overflow

28. The sequence $3D800000$ represents the following number in IEEE 754 single precision floating point representation:

(a) 0.625
(b) 16.25
(c) 41
(d) 221

29. The number 725526135.5202 in octal representation is represented in hexadecimal as:

(a) $7A21FF.031$,
(b) $BES234.FA3$
(c) $ABB783.921$
(d) $EAB58B.A82$

30. Moving from the instruction set architecture level to the microarchitecture level is done by:

(a) translation
(b) interpretation
(c) direct execution

2 Operating Systems

1. Mark the correct answers. An operating system may be seen as...
   (a) An extension of the computing system
   (b) A resource manager
   (c) A control application
   (d) An application program
   (e) A driver for attached devices
   (f) A user graphical interface
2. Mark the correct answers. The functions of an operating system include...
   (a) Offering a convenient interface with the user
   (b) Offering a system for resource management
   (c) Supporting newly identified requirements
   (d) File formatting
   (e) Interpreting command lines
   (f) Compiling user applications
3. Mark the correct answers. The purposes of the resident monitor include the following...
   (a) Grouping of jobs with similar needs
   (b) Automatic succession of jobs
   (c) Interpreting of the punched cards
   (d) Command line interpreting
   (e) Assuring the necessary means for off-line processing
   (f) Assuring the necessary means for multi-programming
4. Mark the correct answers. Multi-programming allows...
   (a) Simultaneous execution of several jobs
   (b) Efficient process (job) scheduling
   (c) Implementation of processor time-sharing
   (d) Implementation of the SPOOLING technique (simultaneous peripheral operations on-line)
   (e) Sequential execution of jobs
   (f) That the operating system is able to make decisions on behalf of users
5. Consider the following operating systems – CP/M, THE, VM/370, Minix. Identify the appropriate structure for each of these operating systems...

(a) CP/M – monolithic; THE – layered; VM/370 – virtual machines; Minix – microkernel
(b) CP/M – layered; THE – virtual machines; VM/370 – microkernel; Minix – client-server
(c) CP/M – virtual machines; THE – microkernel; VM/370 – client-server; Minix – monolithic
(d) CP/M – microkernel; THE – client-server; VM/370 – monolithic; Minix – layered
(e) CP/M – client-server; THE – monolithic; VM/370 – layered; Minix – exo-kernel
(f) CP/M – monolithic; THE – layered; VM/370 – exo-kernel; Minix – virtual machines

6. Simultaneous peripheral operations on-line (spooling) offers support for ...

(a) Intensive use of disk units
(b) Simultaneous use of input/output devices
(c) The introduction of scheduling for the central processing unit
(d) Remote data processing
(e) Time-sharing implementation
(f) The use of the resident monitor

7. Mark the correct answers for threads...

(a) In the hybrid mode, threads in user space are mapped onto operating system's processes
(b) The multithreading is solved in a similar manner to multiprogramming
(c) Applications have the ability to specify points where cancellation of threads is possible
(d) A detached thread can be the subject of a thread-join call
(e) A detached thread can become joinable, and vice versa
(f) The protection mechanisms between threads require a strong support from the operating system

8. Mark the correct answers for process management...

(a) involves providing of synchronization mechanisms between processes
(b) involves providing of communication mechanisms between processes
(c) involves providing a buffering system
(d) requires a decision on the processes to be loaded into main memory
(e) involves saving programs in files
(f) involves ensuring the independence of the execution of processes

9. Consider the model with five states. Which of the following transitions are in relation with the short-term scheduler?

(a) Running process $\rightarrow$ Ready process

(b) Ready process → Running process

(c) Blocked process → Blocked/suspended into memory process

(d) Running process → Blocked process

(e) Ready/suspended into memory process → Ready process

(f) Blocked process → Ready process

10. Consider the model with five states. Which of the following transitions are in relation with the mid-term scheduler?

(a) Running process → Ready process

(b) Ready process → Running process

(c) Blocked process → Blocked/suspended into memory process

(d) Running process → Blocked process

(e) Ready/suspended into memory process → Ready process

(f) Blocked process → Ready process

11. There is a set of four conditions for a good solution of the critical section problem. Which of these were specified in the following list?

(a) No two processes may be simultaneously inside their critical sections.

(b) No assumptions may be made about the speed or the number of processors.

(c) No process running outside its critical section may block other processes.

(d) No process should have to wait forever to enter its critical section.

(e) A process running inside its critical section may block other processes which are outside their critical region.

(f) If no process is running inside its critical region, access to the critical region must be granted to the first process who previously left its critical section.

12. Which of the following IPC mechanisms are based on busy-waiting?

(a) The `lock` variable

(b) Peterson's solution

(c) The TSL instruction

(d) The sleep-wakeup mechanism

(e) Semaphores

(f) The mutex mechanism

13. Which of the following IPC mechanisms involves process suspension?

(a) The sleep-wakeup mechanism

(b) Semaphores
(c) The mutex mechanism
(d) The `lock` variable
(e) Peterson's solution
(f) The TSL instruction

14. Which of the following scheduling algorithms are non-preemptive?

(a) First-In First-Out
(b) Shortest Job First
(c) Shortest Return Time Next
(d) Round Robin
(e) Guarantee Scheduling
(f) Lottery Scheduling

15. Which of the following scheduling algorithms are preemptive?

(a) Shortest Return Time Next
(b) Round Robin
(c) Guarantee Scheduling
(d) Lottery Scheduling
(e) First-In First-Out
(f) Shortest Job First

16. Which of the following scheduling algorithms can degenerate into non-preemptive algorithms?

(a) Shortest Return Time Next → Shortest Job First
(b) Round Robin → First-In First-Out
(c) First-In First-Out → Shortest Job First
(d) Shortest Job First → First-In First-Out
(e) Guarantee Scheduling → Shortest Return Time Net
(f) Lottery Scheduling → Priority Scheduling

17. Which of the following information represents basic requirements for scheduling?

(a) Fairness: offer similar treatment for comparable processes
(b) Equilibrium: ensuring a high level of resource utilization
(c) Policy enforcement: seeing that stated (local) policies are carried out
(d) Output: maximizing the time spent by each process in the computing system

(e) Proportionality: ensuring that the response time is proportional with estimated execution time

(f) Processor utilization: maximize the processor time used by each process

18. A set of four necessary conditions for deadlocks exist. Which of these were specified in the following list?

(a) Mutual exclusion

(b) Hold and wait

(c) Non-preemptive resources

(d) Circular wait

(e) Independence of processes

(f) Preemptive resources

19. Process termination could occur...

(a) After an exit system call

(b) As a result of a program exception

(c) After the eviction of a process from memory to disk

(d) After the execution of any system call

(e) After reading of a control card

(f) After the creation of a child process

20. Process creation could happen...

(a) As a result of a (specific) user request

(b) As a result of a specific system call

(c) On system startup

(d) After the execution of any system call

(e) After reading of a control card, in a batch processing system

(f) On termination of another process

21. Mark the correct answers for main memory...

(a) The main memory loses its content without system power

(b) The main memory stores programs and data for a computing system

(c) The main memory loses its content when another program is loaded

(d) The main memory can be directly accessed by the processor

(e) The main memory offers access to operating system data

(f) The main memory is an intermediary for input-output operations

22. Mark the correct answers for interrupt handling...

(a) In an ordinary system, while processing an interrupt one can handle other interrupts
(b) When an interrupt occurs, the hardware transfers the control to the operating system
(c) The interrupt vector contains addresses of interrupt handling routines
(d) One can handle interrupts in user mode
(e) Events are signaled to the operating systems via software interrupts
(f) In an ordinary system, while processing an interrupt any other interrupts are disabled

23. Mark the correct answers for the protection of instructions...

(a) Privileged instructions can be executed in kernel (monitor) mode
(b) It is possible to execute privileged instructions in user mode
(c) A privileged instruction in user mode determines an interrupt to the operating system
(d) The command line interpretor is offering its services only in monitor mode
(e) The operating system is loaded only in user mode
(f) The mode bit allows users to specify the processes that can be executed in monitor mode

24. Mark the correct answers for security...

(a) The operating system is responsible for carrying out input-output operations
(b) The mode bit allows users to specify the processes that can be executed in monitor mode
(c) In an operating system, users are able to carry out input-output operations directly
(d) The operating system can access only the monitor memory
(e) The base and limit registers can be modified by user applications
(f) The base and limit registers can be loaded by privileged instructions

25. Consider the following set of processes, described by their burst times (3, 6, 12, 4, 15, 8, 9, 5). These processes arrive simultaneously in the system at moment 0 (first four processes) and 15 (last four processes). If the scheduling algorithm is FIFO, mark the correct answers...

(a) The response time for process 3 is 21
(b) The waiting time for process 6 is 25
(c) The number of context switch operations is 7
(d) The response time for process 4 is 21
(e) The waiting time for process 5 is 25
(f) The number of context switch operations is 8

26. Consider the following set of processes, described by their burst times (3, 6, 12, 4, 15, 8, 9, 5). These processes arrive simultaneously in the system at moment 0 (first four processes) and 15 (last four processes). If the scheduling algorithm is SRTN, mark the correct answers...

    (a) The response time for process 3 is 47

    (b) The waiting time for process 6 is 5

    (c) The number of context switch operations is 8

    (d) The response time for process 4 is 3

    (e) The waiting time for process 5 is 47

    (f) The number of context switch operations is 7

27. Consider a simple situation with 6 processes and one resource type. The state of this system is described by $Has = (4, 2, 0, 5, 1, 4), Max = (8, 10, 10, 25, 25, 30), Available = 14$. Use bankers' algorithm and answer the following questions

    (a) This state is sure

    (b) This state is not sure

    (c) After an allocation of 3 resources for the last process, this state remains sure

    (d) After an allocation of 3 resources for the last process, this state is unsure

    (e) After an allocation of 3 resources for the fourth process this state remains sure

    (f) After an allocation of 3 resources for the fourth process, this state is unsure

28. Consider a simple situation with 6 processes and one resource type. The state of this system is described by $Has = (4, 2, 0, 5, 1, 4), Max = (14, 6, 8, 29, 27, 24), Available = 14$. Use bankers' algorithm and answer the following questions

    (a) This state is sure

    (b) This state is not sure

    (c) After an allocation of 3 resources for the last process, this state remains sure

    (d) After an allocation of 3 resources for the last process, this state is unsure

    (e) After an allocation of 3 resources for the fifth process this state remains sure

    (f) After an allocation of 3 resources for the fifth process, this state is unsure

29. In a computing system, the size of a page is $4k$, and 4 frame pages and 8 virtual pages are supported. The following accesses to pages are specified, until the moment 16 (additional information is offered for the optimal algorithm): 0 4 2 4 1 7 3 4 6 5 2 1 6 7 5 2 3 4 5 6 7 0 1 2.

    (a) At moment 12, and the FIFO algorithm, the following pages are mapped in memory: 5216

    (b) For the FIFO algorithm it will be possible to access an address form the virtual page which is mapped on first frame page at moment 7, at moment 16, too.

(c) The number of page faults for the two algorithms are $(9, 12)$, in the order Optimal, FIFO

(d) At moment 12, and the FIFO algorithm, the following pages are mapped in memory: 5217

(e) For the Optimal algorithm it will be possible to access an address form the virtual page which is mapped on third frame page at moment 7, at moment 16, too.

(f) The number of page faults for the two algorithms are $(9, 11)$,in the order Optimal, FIFO

30. In a computing system, the size of a page is $4k$, and 4 frame pages and 8 virtual pages are supported. The following accesses to pages are specified, until the moment 16 (additional information is offered for the optimal algorithm): 0 4 2 4 7 1 4 3 5 6 2 1 6 7 5 2 3 4 5 6 7 0 1 2.

(a) At moment 12, and the Second Chance algorithm, the following pages are mapped in memory: 2 6 1 5

(b) For the Second Chance algorithm it will be possible to access an address form the virtual page which is mapped on first frame page at moment 7, at moment 16, too.

(c) The number of page faults for the two algorithms are $(9, 12)$, in the order Optimal, Second Chance

(d) At moment 12, and the Second Chance algorithm, the following pages are mapped in memory: 5127

(e) For the Optimal algorithm it will be possible to access an address form the virtual page which is mapped on third frame page at moment 7, at moment 16, too.

(f) The number of page faults for the two algorithms are $(9, 11)$, in the order Optimal, Second Chance

3 Databases

1. Which of the following are typical languages used in a DBMS?
   (a) Data Definition Language
   (b) Data Characterization Language
   (c) Data Manipulation Language
   (d) Data Mirroring Language
   (e) Data Modularization Language
2. Which are the components of the three-level architecture?
   (a) Outside level
   (b) Language level
   (c) Conceptual level
   (d) Internal level
   (e) External level
3. Which of the following are properties of a candidate key?
   (a) Uniqueness
   (b) Irreducibility
   (c) Non-uniqueness
4. Which of the following are fundamental operations in relational algebra?
   (a) Selection
   (b) Projection
   (c) Cartesian Product
   (d) Union
   (e) Set Difference
   (f) Set Intersection
5. Which of the following are types of join operations?
   (a) Theta join
   (b) Equi join
   (c) Natural join
   (d) Outer join
   (e) Semi join

6. Which of the following is a correct definition for the third normal form (3NF)?
   (a) A relation that is in first and second normal form, and in which no non-primary key attribute is transitively dependent on the primary key
   (b) A relation that is in first and second normal form, and in which all primary key attributes are transitively dependent on the primary key
   (c) A relation that is in first and second normal form, and in which no non-primary key attribute is functionally dependent on the primary key
7. Which of the following things is the internal level concerned with?
   (a) Storage space allocation for data and indexes
   (b) Record descriptions for storage
   (c) Record placement
   (d) Data compression and data encryption techniques
8. What is a superkey?
   (a) An attribute or a set of attributes that uniquely identifies a tuple within a relation
   (b) An attribute or a set of attributes that uniquely identifies a relation in a database
   (c) An attribute or a set of attributes that is more powerful than simple keys
9. Which of the following represents the entity integrity rule of the relational model?
   (a) Some of the attributes in the primary key can be null
   (b) If an attribute in the primary key is null, it should be ignored
   (c) No attribute of a primary key can be null
10. Which of the following sentences are true about attribute domains?
    (a) Domains determines which classes of entities can be compared
    (b) Domains defines the valid operations for an attribute.
    (c) Domains defines the allowed values for an attribute
11. Which of the following are issues addressed by normalization?
    (a) Deletion anomalies
    (b) Insertion anomalies
    (c) Modification anomalies
    (d) Data redundancy
    (e) Data volume

12. Consider the relation R(A, B, C, D, E) with the following functional dependencies: D $\rightarrow$ C, {C, E} $\rightarrow$ A, D $\rightarrow$ A, {A, E} $\rightarrow$ D. Which of the following attribute sets is a candidate key in R?

    (a) {A}

    (b) {B, D, E}

    (c) {C, D, E}

    (d) {A, D}

13. Three of the following four expressions finds the names of all students who did not apply to major in CS or PH. Which one finds something different? (Hint: You should not assume student names are unique.)

    (a) $\pi_{sName}(Student \bowtie$
    $(\pi_{sID}(Student) - (\pi_{sID}(\sigma_{major='CS'}Apply) \cup \pi_{sID}(\sigma_{major='PH'}Apply))))$

    (b) $\pi_{sName}(Student \bowtie$
    $(\pi_{sID}Student - \pi_{sID}(\sigma_{major='CS' \vee major='PH'}Apply)))$

    (c) $\pi_{sName}(\pi_{sID,sName}Student -$
    $\pi_{sID,sName}(Student \bowtie \pi_{sID}(\sigma_{major='CS' \vee major='PH'}Apply)))$

    (d) $\pi_{sName}Student -$
    $\pi_{sName}(Student \bowtie \pi_{sID}(\sigma_{major='CS' \vee major='PH'}Apply))$

14. Suppose relation Student has 20 tuples. What is the minimum and maximum number of tuples in the result of this expression

    $\rho_{s1(i1,n1,g,h)}(Student) \bowtie \rho_{s2(i2,n2,g,h)}(Student)$

    (a) minimum = 0, maximum = 400

    (b) minimum = 20, maximum = 400

    (c) minimum = 40, maximum = 40

    (d) minimum = 20, maximum = 20

15. Consider the following SQL table declaration:

```
CREATE TABLE Employees (id INT,
                        ssNo INT,
                        name CHAR(20),
                        managerID INT);
```

    We would like to extend the table declaration to enforce that each of id and ssNo is a key (by itself), and each value of managerID must be one of the values that appears in the id attribute of the same table. Which of the following SQL additions will accomplish these?

    (a) Add "UNIQUE" after the first INT, and add "PRIMARY KEY" after the second INT.

(b) Add ", CONSTRAINT PK PRIMARY KEY (id, ssNo)" before the closing parenthesis.

(c) Add "PRIMARY KEY" after the first INT, and add "UNIQUE" after the last INT.

(d) Add "REFERENCES Employees(id)" before the closing parenthesis.

16. Here are SQL declarations for two tables S and T:

```
CREATE TABLE S(c INT PRIMARY KEY, d INT);
CREATE TABLE T(a INT PRIMARY KEY, b INT REFERENCES S(c));
```

Consider S(c, d) contains the tuples {(2, 10), (3, 11), (4, 12), (5, 13)} and T(a, b) contains {(0, 4), (1, 5), (2, 4), (3, 5)}. As a result of the constraints in the table declarations, certain insertions, deletions, and/or updates on S and T are disallowed. Which of the following modifications will not violate any constraint?

(a) Inserting (6, 1) into T

(b) Deleting (2, 4) from T

(c) Deleting (5, 13) from S

(d) Deleting (4, 12) from S

(e) Deleting (3, 11) from S

17. Given a relation R(A) containing the tuples {(1), (2)} and two transactions

```
(T1) UPDATE R SET A = 2*A
(T2) SELECT AVG(A) FROM R
```

If transaction T2 executes using *read committed*, what are the possible values it returns?

(a) 1

(b) 1.5

(c) 2

(d) 2.5

(e) 3

18. Consider the relations R(A,B) with two tuples {(1, 5), (2, 5)} and S(B,C) with one tuple {(5, 10)}, and a view V defined by the following query

```
SELECT A, C FROM R, S WHERE R.B = S.B
```

Which of the following statements is/are true?

(a) after deleting (2, 5) from R the cardinality of V becomes 1

(b) updating (2, 5) to (2, 6) in R will leave V empty

(c) updating (2, 5) to (3, 5) in R will leave V unchanged

(d) after deleting (5, 10) from S the cardinality of V becomes 0

(e) updating (1, 5) to (2, 6) in R will delete the tuple (2, 10) from V

19. Consider the following trigger over a relation R(a, b):

```
CREATE TRIGGER InsertOnR
AFTER INSERT ON R
REFERENCING NEW ROW AS new
FOR EACH ROW
WHEN (new.a * new.b > 30)
INSERT INTO R VALUES (new.a - 1, new.b + 1);
```

Suppose we begin with table R empty. Which of the insertions below results in R containing exactly 5 tuples?

(a) (50, 0)

(b) (5, 3)

(c) (7, 5)

(d) (11, 1)

20. Consider the following query:

```
SELECT * FROM Student, Apply, College
WHERE Student.sID = Apply.sID
  AND Apply.cName = College.cName
  AND Student.Bac > 5
  AND College.cName = 'UVT'
```

Suppose we can create two indexes, and assume all indexes are tree-based. Which two indexes do you think would be most useful for speeding up query execution?

(a) Student.sID, Student.Bac

(b) Apply.cName, College.cName

(c) Apply.sID, College.cName

(d) Apply.sID, Student.Bac

21. Consider the relation R(A, B, C, D, E, F) with functional dependencies A $\rightarrow$ {B, C, E} and E $\rightarrow$ F. This relation is in which normal form (NF)?

(a) 1NF

(b) 2NF

(c) 3NF

(d) BCNF

(e) 4NF

22. Evaluate the following statements:

```
CREATE TABLE digits(id NUMBER(2), description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO' WHERE id=1;
INSERT INTO digits VALUES (2,'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

What would be the outcome of the above query?

(a) It would not display any values.

(b) It would display the value TWO once.

(c) It would display the value TWO twice.

(d) It would display the values ONE, TWO, and TWO.

23. You want to know the FIRST_NAME and SALARY for all employees who have the same manager as that of the employee with the first name 'Neena' and have salary equal to or greater than that of 'Neena'.
Which SQL statement would give you the desired result?

(a) SELECT first_name, salary FROM employees WHERE (manager_id, salary) >= ALL (SELECT manager_id, salary FROM employees WHERE first_name = 'Neena') AND first_name <> 'Neena';

(b) SELECT first_name, salary FROM employees WHERE (manager_id, salary) >= (SELECT manager_id, salary FROM employees WHERE first_name = 'Neena') AND firs_name <> 'Neena';

(c) SELECT first_name, salary FROM employees WHERE (manager_id, salary) >= ANY (SELECT manager_id, salary FROM employees WHERE first_name = 'Neena' AND first_name <> 'Neena';

(d) SELECT first_name, salary FROM employees WHERE (manager_id = (SELECT manager_id FROM employees WHERE first_name = 'Neena') AND salary >= (SELECT salary FROM employees WHERE first_name = 'Neena')) AND first name <> 'Neena';

24. EMPDET is an external table containing the columns EMPNO and ENAME. Which command would work in relation to the EMPDET table?

(a) UPDATE empdet SET ename = 'Amit' WHERE empno = 1234;

(b) DELETE FROM empdet WHERE ename LIKE 'J%';

(c) CREATE VIEW empvu AS SELECT * FROM empdet;

(d) CREATE INDEX empdet_idx ON empdet(empno);

25. Evaluate the CREATE TABLE statement:

```
CREATE TABLE products (
    product_id NUMBER(6) CONSTRAINT prod_id_pk PRIMARY KEY,
    product_name VARCHAR2(15));
```

Which statement is true regarding the PROD_ID_PK constraint?

(a) It would be created only if a unique index is manually created first.

(b) It would be created and would use an automatically created unique index.

(c) It would be created and would use an automatically created non-unique index.

(d) It would be created and remains in a disabled state because no index is specified in the command.

26. Evaluate the following SQL query:

```
SELECT TRUNC(ROUND(159.99, 1),1) FROM DUAL;
```

What would be the outcome?

(a) 16

(b) 100

(c) 160

(d) 200

(e) 150

27. You need to calculate the number of days from 1st Jan 2007 till date: Dates are stored in the default format of dd-mm-yy. Which two SQL statements would give the required output? (Choose two.)

(a) SELECT SYSDATE - '01-JAN-2007' FROM DUAL

(b) SELECT SYSDATE - TO_DATE('01/JANUARY/2007') FROM DUAL;

(c) SELECT SYSDATE - TO_DATE('01-JANUARY-2007') FROM DUAL;

(d) SELECT TO_CHAR(SYSDATE,'DD-MON-YYYY')-'01-JAN-2007' FROM DUAL;

(e) SELECT TO_DATE(SYSDATE,'DD/MONTH/YYYY')-'01/JANUARY/2007' FROM DUAL;

28. You want to display all the cities that have no departments and the departments that have not been allocated cities. Which type of join between DEPARTMENTS and LOCATIONS tables would produce this information as part of its output?

    (a) NATURAL JOIN

    (b) FULL OUTER JOIN

    (c) LEFT OUTER JOIN

    (d) RIGHT OUTER JOIN

29. Evaluate the following statement:

```
INSERT ALL
WHEN order_total < 10000 THEN
INTO small_orders
WHEN order_total > 10000 AND order_total < 20000 THEN
INTO medium_orders
WHEN order_total > 2000000 THEN
INTO large_orders
SELECT order_id, order_total, customer_id FROM orders;
```

    Which statement is true regarding the evaluation of rows returned by the subquery in the INSERT statement?

    (a) They are evaluated by all the three WHEN clauses regardless of the results of the evaluation of any other WHEN clause.

    (b) They are evaluated by the first WHEN clause. If the condition is true, then the row would be evaluated by the subsequent WHEN clauses.

    (c) They are evaluated by the first WHEN clause. If the condition is false, then the row would be evaluated by the subsequent WHEN clauses.

    (d) The INSERT statement would give an error because the ELSE clause is not present for support in case none of the WHEN clauses are true.

30. What does the following SQL query do?

```
SELECT b.Name
FROM Employee a, Employee b, SalaryLevel c, SalaryLevel d
WHERE a.ManagerId = b.Id
  AND b.Salary BETWEEN c.MinSalary AND c.MaxSalary
  AND b.Salary*1.25 BETWEEN d.MinSalary AND d.MaxSalary
  AND c.Level+1 = d.Level
```

    (a) Returns employees' names who would move to the next salary level after a 25% wage increase.

(b) Returns managers' names who would move to the next salary level after a 25% wage increase.

(c) Returns the employees who would move to the next salary level after a 25% wage increase if their salary is between the minimum and maximum wage.

4 Computer Networks

1. The physical layout of a computer network is known as:

   (a) protocol;
   (b) topology;
   (c) backbone;
   (d) segment;

2. What is the network topology where each node is connected to the two nearest nodes so that the entire network forms a circle?

   (a) bus;
   (b) ring;
   (c) star;
   (d) bus-star;

3. What are some of the disadvantages of a ring topology?

   (a) it does not scale well;
   (b) a single malfunctioning workstation can disable the entire network;
   (c) it is very flexible;
   (d) it is a bidirectional topology;

4. A process socket local address is equal to:

   (a) port number + IP address;
   (b) IP address;
   (c) port number;
   (d) IP address + hostname + port number;

5. To ensure data integrity further, connection-oriented protocols such as TCP use a:

   (a) digital signature;
   (b) digital certificate;
   (c) symmetric encryption algorithm;
   (d) checksum;

6. What kind of transport layer protocols are more useful in situations where data must be transferred quickly?

   (a) connectionless protocols;

(b) syn-oriented protocols;
(c) connection-oriented protocols;
(d) ack-oriented protocols;

7. PDUs at the OSI network layer are called what?

(a) transport;
(b) frames;
(c) packets;
(d) segments;

8. Network layer protocol that reports on the success or failure of data delivery:

(a) IP;
(b) TCP;
(c) ARP;
(d) ICMP;

9. When data is encapsulated, which is the correct order?

(a) data, frame, packet, segment, bit;
(b) segment, data, packet, frame, bit;
(c) data, segment, packet, frame, bit;
(d) data, segment, frame, packet, bit;

10. Acknowledgments, sequencing, and flow control are characteristic of which OSI layer?

(a) layer 2;
(b) layer 3;
(c) layer 4;
(d) layer 7;

11. On what layer of the OSI model does TCP operate?

(a) physical;
(b) data link;
(c) session;
(d) transport;

12. Which of the following services use UDP?

(a) SMTP;

(b) SNMP;
(c) FTP;
(d) TFTP;
(e) DHCP;
(f) HTTP;

13. What can be said about TCP?

(a) it is a connectionless protocol;
(b) it is a connection-oriented protocol;
(c) it does not use checksums;
(d) it provides segmentation and reassembly;

14. What is the name of the protocol that allows a client to send a broadcast message with its MAC address and receive an IP address in reply?

(a) ARP;
(b) DNS;
(c) RARP;
(d) ICMP;

15. What type of cable consists of a copper core, metal shielding, and a jacket?

(a) fiber-optic;
(b) UTP;
(c) twisted-pair;
(d) coaxial;

16. What IP address will you use to send a message to all devices connected to your network segment?

(a) 0.0.0.0;
(b) 127.0.0.1;
(c) 255.0.0.0;
(d) 255.255.255.255;

17. The IP address 127.0.0.1 is also known as:

(a) loopback address;
(b) broadcast address;
(c) multicast address;
(d) class A broadcast address;

18. Which of the following devices are layer 1 devices?

    (a) bridge;
    (b) repeater;
    (c) router;
    (d) switch;
    (e) hub;

19. Which of the following devices are layer 2 devices?

    (a) bridge;
    (b) repeater;
    (c) router;
    (d) switch;
    (e) hub;

20. Which of the following devices are layer 3 devices?

    (a) bridge;
    (b) repeater;
    (c) router;
    (d) switch;
    (e) hub;

21. Which of the following types of connections can use full-duplex?

    (a) hub to hub;
    (b) switch to switch;
    (c) host to host;
    (d) switch to hub;
    (e) switch to host;

22. What protocol is used to find the hardware address of a local device?

    (a) RARP
    (b) ARP
    (c) IP
    (d) ICPM
    (e) BootP

23. What is the result of segmenting a network with a bridge?

    (a) it increases the number of collision domains;
    (b) it decreases the number of collision domains;
    (c) it increases the number of broadcast domains;
    (d) it decreases the number of broadcast domains;
    (e) it makes smaller collision domains;
    (f) it makes larger collision domains;

24. What are the main functions of layer 2 switches?

    (a) address learning;
    (b) routing;
    (c) forwarding and filtering;
    (d) creating network loops;
    (e) loop avoidance;
    (f) IP addressing;

25. Which of the following are reasons for breaking up a network into two segments with a router?

    (a) to create fewer broadcast domains;
    (b) to create more broadcast domains;
    (c) to create one large broadcast domain;
    (d) to create one large collision domain;

26. Which of the following are benefits in using fiber cabling?

    (a) reliability;
    (b) requires special equipment;
    (c) throughput;
    (d) low security levels;

27. What are the TCP/IP protocols that run at the transport layer of the OSI model and provides reliable data delivery services?

    (a) UDP
    (b) IP
    (c) TCP
    (d) ARP
    (e) HTTP

28. On a TCP segment, what field allows the receiving node to determine whether the TCP segment became corrupted during transmission?

    (a) checksum;

    (b) flags;

    (c) hash;

    (d) padding;

29. You need 500 subnets, each with about 100 usable host addresses per subnet. What mask will you assign using a Class B network address?

    (a) 255.255.255.252;

    (b) 255.255.255.128;

    (c) 255.255.255.0;

    (d) 255.255.254.0;

30. Which of the following IP addresses fall into the CIDR block of 115.64.4.0/22? (Choose two.)

    (a) 115.64.8.32;

    (b) 115.64.6.255;

    (c) 115.64.8.32;

    (d) 115.64.5.128;

5 Answers

Computer Architecture

1. 1a,1c
2. 2a
3. 3b
4. 4a,4b,4c
5. 5c
6. 6b
7. 7b
8. 8a,8b,8c
9. 9a,9b
10. 10c
11. 11a
12. 12b,12c
13. 13b,13c
14. 14b
15. 15c
16. 16c
17. 17a
18. 18a,18b
19. 19b
20. 20b
21. 21a
22. 22a,22c,22d,22e
23. 23b
24. 24a,24c
25. 25a,25c
26. 26a
27. 27a
28. 28a
29. 29d
30. 30b

Operating Systems

1. 1a,1b,1c
2. 2a,2b,2c
3. 3a,3b,3c
4. 4a,4b,4c
5. 5a
6. 6a,6b,6c
7. 7a,7b,7c
8. 8a,8b
9. 9a,9b
10. 10c,10d
11. 11a,11b,11c,11d
12. 12a,12b,12c
13. 13a,13b,13c
14. 14a,14b
15. 15a,15b,15c,15d
16. 16a,16b
17. 17a,17b,17c
18. 18a,18b,18c,18d
19. 19a,19b
20. 20a,20b
21. 21a,21b
22. 22a,22b
23. 23a,23b
24. 24a,24b
25. 25a,25b,25c
26. 26a,26b,26c
27. 27a,27d,27e
28. 28a,28c,28f
29. 29a,29b,29c
30. 30a,30b,30c

Databases

1. 1a,1c
2. 2c,2d,2e
3. 3a,3b
4. 4a,4b,4c,4d,4e
5. 5a,5b,5c,5d,5e
6. 6a
7. 7a,7b,7c,7d
8. 8a
9. 9c
10. 10b,10c
11. 11a,11b,11c,11d
12. 12b
13. 13d
14. 14b
15. 15a, 15d
16. 16b, 16e
17. 17b, 17e
18. 18a, 18d
19. 19c
20. 20c
21. 21a
22. 22c
23. 23d
24. 24c
25. 25b
26. 26c
27. 27b,27c
28. 28b
29. 29a
30. 30b

Computer Networks

1. 1b
2. 2b
3. 3a,3b
4. 4a
5. 5d
6. 6a
7. 7c
8. 8d
9. 9c
10. 10c
11. 11d
12. 12b,12d,12e
13. 13b,13d
14. 14c
15. 15d
16. 16d
17. 17a
18. 18b,18e
19. 19a,19d
20. 20c
21. 21b,21c,21e
22. 22b
23. 23a,23e
24. 24a,24c,24e
25. 25b
26. 26a,26c
27. 27c
28. 28a
29. 29b
30. 30b,30d