Entrance exam - Computer Science

Name and Surname – fill in the field	*Application No.*	*Test Sheet No.*
		2

Algorithms and Data Structures

1 Which statement is true?

A The various kinds of binary search trees include AVL trees and binary heaps.
B The input condition of an algorithm describes those inputs for which the algorithm's time complexity is optimal.
***C** The function *n * log(n)* grows asymptotically as fast as the function *log(n!)*.
D The hash table data structure is well-suited for sorting sequences of numbers.
E The time complexity of inserting an element into a binary search tree is always *O(log n)*.

---

2 Consider a directed graph that contains a path from the vertex *u* to the vertex *v*. Which statement is true?

A None of the other statements is true.
B In the breadth-first search algorithm, the vertex *u* is discovered earlier than the vertex *v*.
C In the depth-first search algorithm, the vertex *u* is discovered earlier than the vertex *v*.
***D** For every vertex *s* such that there is a path from *s* to *u*, there is also a path from *s* to *v*.
E For every vertex *s* such that there is a path from *s* to *v*, there is also a path from *s* to *u*.

---

3 Consider a binary min heap with *n* elements. Which statement is false?

A The time complexity of removing the minimum element of the heap is *O(log n)*.
B The time complexity of finding the minimum element of the heap is *O(1)*.
***C** The time complexity of removing the minimum element of the heap is *O(1)*.
D The time complexity of finding a given element in the heap is *O(n)*.
E The time complexity of finding the minimum element of the heap is *O(log n)*.

---

4 Consider a hash table whose collision resolution is based on separate chaining with linked lists. The size of the hash table is 7. The hash function used is *h(x) = (3*x + 5) mod 7*. We start with an empty hash table and subsequently insert the following values: 12, 2, 5, 7, 9, 1. Which statement about the state of the hash table is true after all the insertions?

***A** None of the other statements is true.
B One of the hash table slots contains a linked list with the elements 7 and 12.
C The longest linked list in the table has a length of 3 elements.
D One of the hash table slots contains a linked list with the elements 2 and 5.
E The longest linked list in the table has a length of 6 elements.

5

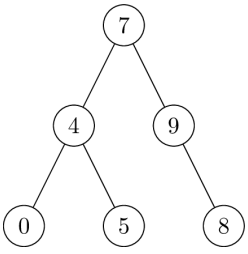

Consider the tree depicted above. Which one of the following operations should we perform to change the tree into a binary search tree?

A Swap the values 7 and 9.
B Swap the values 7 and 8.
***C** Remove the vertex with the value 8.
D Move the value 9 into the root, the value 7 into the rightmost leaf, and the value 8 into the rightmost vertex of the second level.
E None, the tree already is a binary search tree.

Computer systems

6 Let the CPU scheduling use the SJF (Shortest Job First) algorithm in the preemptive variant. The request for process P1 is coming in time 0 and the process needs 6 units of CPU. The request for process P2 is coming in time 2 and the process needs 3 units of CPU. The request for process P3 is coming in time 2 and the process needs 2 units of CPU. The request for process P4 is coming in time 3 and the process needs 2 units of CPU. The CPU scheduling starts in time 0. Which process will be running at the CPU in time 7?

A none of these processes
***B** P2
C P3
D P1
E P4

---

7 Which binary number is equivalent to the decimal number 23456?

A 0111 1010 1010 1000
B 55640
C 0000 1010 1011 0101
***D** 0101 1011 1010 0000
E 1010 0100 0101 1111

---

8 The longterm (strategic) scheduler in an operating system:

A selects one of ready processes to run on the CPU.
***B** decides which program (job) and when it will be executed.
C changes the status of the thread/process from waiting/blocked to ready.
D must be fast, because its code is executed very frequently.
E decides which executed process will be swapped out to disk.

---

9 A necessary condition of a deadlock is:

A the need to use at least three different resources.
B existence of at least two instances of each resource.
C using a semaphore.
D possibility of concurrent use of a single instance of a resource by at least two threads.
***E** keeping an allocated resource and waiting for another one (hold and wait).

---

10 Each thread in a process has:

A its own address space which can be shared with other threads using IPC mechanisms.
B its own table of open files.
C its own file locks.
D its own signal handlers.
***E** its own stack.

Programming

11 Which statement is generally true in common OOP languages (C++, Java, C#):

A The difference between a class and an object is that objects can inherit from classes but not the other way round.
***B** If a class B inherits from a class A, B can access all attributes of A that are declared public or protected.
C None of the other statements is true.
D If early binding (non-virtual method calls) is used, the actual method to be called is decided at run time.
E If a method of a class is declared as static, it means that it can access private attributes of the class; non-static methods cannot access those.

---

12 Which of the following three statements I, II, and III are true (in common OOP languages such as C++, Java, C#). Choose the option that contains exactly all the true statements (and none of the false ones).

I. Local variables of functions are allocated on the stack. They are automatically deallocated when leaving the function.

II. The memory that is allocated using the operator `new` is allocated on the heap.

III. If an exception is not caught inside the function that threw the exception, it always results in the termination of the whole process.

***A** I, II
B I, II, III
C I, III
D I
E II, III

---

13 Which statement is **false**?

A Prolog is an example of a logic programming language.
***B** If call-by-value is used, the update to a function's parameter is also propagated outside.
C The lazy evaluation strategy in functional programming allows working with infinite lists.
D Recursive function calls are typically implemented using a stack.
E If call-by-reference is used, the update to a function's parameter is also propagated outside.

14

```
function foo(integer n)
begin
        print n
        if n > 0 then
                print "*"
                foo(n-1)
        end if
        print n
end

program main()
begin
        foo(3)
end
```

Assume that the function `print` does not output end of line. What is the output of the program?

A `3*2*1*00`
***B** `3*2*1*00123`
C `3*2*1`
D `3*2*1*123`
E The program will run forever and never halt.

15
```
read(n)
sum = 0
k = n
XXX {
        sum = sum + (n mod 10)
        n = n div 10
}
print sum
```
Let `mod` be the modulo operator and let `div` be the integer division operator. Assume that the input `n` is a positive integer. Which of the following should we place instead of "XXX" in the above code, so that it computes the sum of the digits of `n`?

A `for i = 1 to k`
***B** `while n > 0`
C `for i = 1 to 10`
D None of the other options is true.
E `while n > sum`

Computer Networks

---

16 Path Vector routing protocols for routing between autonomous systems:

A require use of the same metrics for all autonomous systems.
***B** are derived from Distance Vector routing protocols, allow definition of routing rules (policies) and prefer the shortest paths, if it is permitted by the rules.
C are derived from Distance Vector routing protocols and prefer always the shortest paths.
D exchange information about paths, but do not allow to detect cycles.
E are derived from Link State routing protocols and allow definition of routing rules (policies).

---

17 The TCP protocol provides these features:

A routing data with quality of service.
B data flow support for real time applications.
C routing data with quality of service, and congestion and flow control.
***D** creation of data blocks with header, their numbering, connectivity, errors and flow control, and congestion control.
E congestion and flow control and error control on active elements.

---

18 Omnidirectional wireless transmission spreads the emitted energy in all directions. The received energy by the receiver:

A decreases exponentially with the increasing distance of the receiver from the transmitter.
B decreases cubically with the increasing distance of the receiver from the transmitter.
C is identical to the emitted energy regardless of the distance of the receiver from the transmitter.
***D** decreases quadratically with the increasing distance of the receiver from the transmitter.
E decreases linearly with the increasing distance of the receiver from the transmitter.

---

19 The highest achievable transmission rate (bits / second) of a communication channel is determined by:

A the method of coding data into a signal in a communication channel.
***B** the width of the frequency interval of component signals representing signal carrying data and the ratio of transmission energy and harmful, unwanted energies present in the transmission medium.
C the method of coding data into a signal in a communication channel and data volumes, which the transmitter and receiver are able to process per unit of time.
D the width of the frequency interval of component signals representing signal carrying data and method of coding data into the signal.
E the width of the frequency interval of component signals representing signal carrying data.

---

20 Ensuring network services by packet switching:

- ***A** is realized by sending independent data units (packets) and is the cornerstone of connection-less networks.
- **B** is realized by a direct physical connection between the sender and receiver.
- **C** is realized by sending independent data units (packets) through established physical connections.
- **D** provides a connection-oriented service without the need of a packetization.
- **E** is a keystone of connection-oriented networks.

Databases

21 Let there be a relation *r(A, B, C, D, E, F, G, H, I)* and the following set of functional dependencies:
A, B, C, D → E, F
A, C, D, E → B
E → G
A, E → H
A, B → I
Which one of the following statements is true for the set of attributes *{A, B, C, G, H, I}*?

- **A** It is a candidate key of the relation but not a super key.
- **B** It is a super key of the relation but not a candidate key.
- **C** It is the one and only primary key of the relation.
- ***D** It is neither a super key nor candidate key of the relation.
- **E** It is both a super key and a candidate key of the relation.

22 Consider a database of a bank that has the following two relations: *customer(cid, name, address)*, and *account(ban, cid, balance)*. Suppose that neither relation is empty and that indexes exist only for the primary keys. Select the only command that **does not accelerate** the processing of the following SQL query:

```
SELECT * FROM customer, account WHERE customer.cid = account.cid AND balance = 0;
```

- **A** CREATE INDEX speedup ON account(balance);
- **B** ALTER TABLE customer DROP address;
- **C** DELETE FROM account WHERE balance <> 0;
- **D** DELETE FROM account WHERE balance = 0;
- ***E** CREATE INDEX speedup ON customer(name);

23 Let there be a relation *employee* with relational schema *employee(ID, name, salary)*. The domain of attributes *ID* and *salary* is a set of integer numbers, the domain of the attribute *name* is a set of strings of length of up to 10 characters. Which one of the following expressions is a correct example of such a relation (in relational algebra)?

- **A** employee = (1,'Jan',1000,2,'Jana',2000,3,'Jitka',3000)
- **B** employee = {{1,'Jan',1000}, {2,'Jana',2000}, {3,'Jitka',3000}}
- **C** employee = {1,'Jan',1000,2,'Jana',2000,3,'Jitka',3000}
- **D** employee = ({1,'Jan',1000}, {2,'Jana',2000}, {3,'Jitka',3000})
- ***E** employee = {(1,'Jan',1000), (2,'Jana',2000), (3,'Jitka',3000)}

24 Let there be the following relations: *customer(cid, name, address)*, *car(vrp, type)*, and *rental(id, cid, vrp, date_from, date_to, price)*. Suppose that there is at least one record in each relation and that there are no NULL values. The attributes *cid* and *vrp* in the relation *rental* are foreign keys to the relations *customer* and *car* respectively.

What is the result of the following SQL query:

```
SELECT name, COUNT(id) FROM customer NATURAL LEFT OUTER JOIN rental GROUP BY cid, name;
```

- **A** The name and the number of all customers that have rented at least one car.
- **B** The name and the number of rented cars of each customer that has rented a car.
- **C** The name and the number of all cars that have been rented to at least one customer.
- **D** The names of all customers and the total number of cars that have been rented at least once.
- ***E** The name of each customer and the number of cars ever rented.

25 The second normal form (2NF) requires that:

A the ACID rule holds for all attributes.
B all attributes that are part of any foreign key are atomic.
C all dependencies are atomic and that the first normal form (1NF) holds.
D there is no attribute that is transitively dependent on a candidate key.
***E** there is no attribute that is dependent on a part of a candidate key.

Software engineering

26 As a software engineer you shall decide whether a newly developed system shall be designed for a single-server deployment, or shall rather be designed as distributed (deployed on multiple servers). Consider the following three arguments:
I. If the performance of the single-server solution is satisfactory, one shall prefer the single-server deployment, because shift to a distributed environment may negatively affect reliability and security of the system.
II. If it is acceptable from the financial point of view, one should always prefer distributed deployment, because it guarantees higher performance, reliability and security.
III. If it is acceptable from the financial point of view, one should always prefer distributed deployment, because it guarantees higher performance, security and maintainability.

Decide, which of the given arguments I., II. and III. are in general valid (choose the option containing all and only the valid arguments):

A I. and II.
B II.
***C** I.
D II. and III.
E III.

27 Which one of these statements best reflects the changes in the class diagram when transiting from the analysis to the design phase (in case of proper analysis and design)?

A Neither the number of classes nor their complexity increases, but the arguments and return values of their methods are refined.
B The number of classes does not increase, but their complexity increases substantially.
C Neither the number of classes nor their complexity increases, but the relationships among classes are refined.
D No changes occur in the class diagram when transiting from the analysis to the design phase.
***E** The number of classes increases substantially, but their complexity does not increase substantially.

28 For the following activities - Test planning, Cost estimation, Requirements specification review - indicate in which phase of the software life cycle you would recommend them to be done:

A System design
B System implementation
C System maintenance
***D** Requirements specification
E System shutdown

29 Software architecture design tends to be motivated by multiple purposes. Which ones of the following belong to them (i.e. are reasonable purposes of architecture design)?
I. To facilitate reuse (of software components).
II. To facilitate communication (between development teams and other stakeholders).
III. To break down system complexity and support its distribution.
IV. To facilitate system maintenance.

A Only II. and IV., not I. and III.
B Only I., II. and III., not IV.
C Only I. and III., not II. and IV.
***D** All I., II., III., IV.
E Only III. and IV., not I. and II.

30 Consider the following examples of relationships in the UML class diagram. Which one of the following gives the best fitting examples of their proper usage?

A Car-Driver is a composition, Car-VolvoXC90 is an inheritance, Car-Engine is a general association.
B Car-Driver is a composition, Car-VolvoXC90 is a general association, Car-Engine is an inheritance.
C Car-Driver is an inheritance, Car-VolvoXC90 is a general association, Car-Engine is a composition.
***D** Car-Driver is a general association, Car-VolvoXC90 is an inheritance, Car-Engine is a composition.
E Car-Driver is a general association, Car-VolvoXC90 is a composition, Car-Engine is an inheritance.