Entrance exam - Computer Science

Jméno a příjmení – pište do okénka	*Číslo přihlášky*	*Číslo zadání*
		1

Algorithms and Data Structures

1 Which one of these statements is true?

A B+ trees are a special case of binary search trees.

B We say that an algorithm is partially correct, if it produces a correct output for at least some of the inputs.

C The worst-case time complexity of inserting an element into a hash table is in $\mathcal{O}(1)$.

D There exists an $\mathcal{O}(\log n)$ algorithm that creates a binary heap from an unsorted array of elements.

*E The functions $n \log n$ and $\log(n!)$ have the same asymptotic growth.

2 Which of the following is always true for any binary search tree (BST)?

A The depth of the tree is at most five times the logarithm of the number of all nodes.

B The minimum key is stored in the root.

C Every internal node (i.e. not a leaf) references exactly two children.

D The minimum key is stored in one of the leaves.

*E The node with the maximum key does not have a right child.

3 Consider a directed graph. We run a depth-first search on this graph; the search assigns to each vertex *v* two numbers: *v.d* is the discovery time of *v*, *v.f* is the finishing time of *v*. Which one of these statements is true in general?

A If there is a path from vertex *u* to vertex *v* in the graph then *u.f* ≥ *v.f*.

B If there is a path from vertex *u* to vertex *v* in the graph then *u.f* ≤ *v.f*.

C If there is a path from vertex *u* to vertex *v* in the graph then *u.d* ≥ *v.d*.

*D None of the other statements is true in general.

E If there is a path from vertex *u* to vertex *v* in the graph then *u.d* ≤ *v.d*.

4 Consider the maximum binary heap structure implemented as an array of integers. We insert the following keys in this order: 2, 8, 12, 1, 5, 7, 9 . Then we extract the maximum value once. Which one of the following options represents contents of the array?

A [2, 5, 7, 8, 9, 12]

*B [9, 5, 8, 1, 2, 7]

C [1, 2, 5, 7, 8, 9]

D [9, 8, 7, 5, 2, 1]

E [9, 7, 8, 2, 5, 1]

5 Let $f(n) \in \mathcal{O}(n)$ and $g(n) \in \mathcal{O}(n^2)$. Which of the following statements is necessarily true?

A $g(n) \in \mathcal{O}(f(n))$

B $f(n)/g(n) \in \mathcal{O}(1)$

C $f(n) \in \mathcal{O}(g(n))$

D $g(n)/f(n) \in \mathcal{O}(1)$

*E $f(n) \cdot g(n) \in \mathcal{O}(n^3)$

Computer Systems

6 The *FLAGS* register of Intel x86 microprocessors does **not** contain:

A Zero flag

B Overflow flag

C Carry flag

*D Control flag

E Parity flag

7 Which of the following terms does **not** represent one of possible process states?

A   Waiting
*B Concurrent
C   Ready
D   Terminated
E   Running

8 Which of the following sequences represent memory hierarchy by response time?

A   Processor cache, registers, main memory, solid-state drive, magnetic disk, optical disk, magnetic tape
*B Registers, processor cache, main memory, solid-state drive, magnetic disk, optical disk, magnetic tape
C   Registers, processor cache, main memory, solid-state drive, optical disk, magnetic tape, magnetic disk
D   Registers, main memory, processor cache, solid-state drive, magnetic disk, optical disk, magnetic tape
E   Registers, processor cache, main memory, solid-state drive, magnetic disk, magnetic tape, optical disk

9 Consider an unsigned decimal number 173 which is stored in an 8-bit register. Which number
will be in the register after performing a left rotation (circular shift) **twice**?

A   107
B   91
*C 182
D   214
E   145

10 Consider the logic circuit given below.
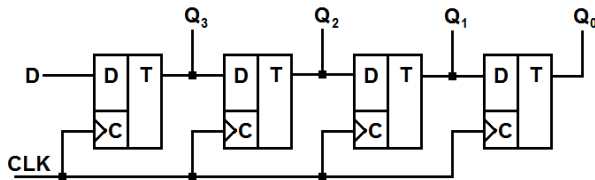
This circuit serves as

A   Parallel register
*B Shift register
C   Multiplexer
D   Counter
E   Frequency divider

Programming

11 Let us consider the following program. The `print` instruction outputs the given number followed by an end-of-line character.

```
function foo(integer n)
begin
        print n
        if n > 0 then
                foo(n-1)
                print n
                foo(n-2)
                print n
        end if
end

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

How many lines are going to be printed by the program?

A   3
B   16
C   29
*D 17
E   The program will run forever and never halt.

12 Which of the following statements is **false**?

A A tail-recursive function can always be rewritten in an iterative manner.
*B When using call-by-value, the change of a function's parameter value by the function's code can be observed from the outside of the function.
C In purely functional languages, functions cannot have any side effect.
D A recursive function can always be rewritten in an iterative manner.
E The normal-order evaluation strategy in functional programming allows working with infinite data structures.

13 Which of the following three statements I, II, and III are true (in common languages such as C++, Java, C#)? Choose the option that contains all the true statements and none of the false ones.
I. Local variables of functions are always allocated on the heap.
II. Function calls are implemented using the stack.
III. If an exception is caught (in a `catch` block), it can be re-thrown (using `throw`).

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

14 Let us consider the following function:
```
function fun(unsigned integer n)
begin
        result = 1
        while n != 0
                n = n - 1
                result = result + 2 * n
        end while
        return result
end
```
Assume that `n > 0`. What is the result computed by `fun`?

A 2^n

B $2n - 1$

C $2n - 2$

*D $n^2 - n + 1$

E n^2

15 Which statement is generally true in common OOP languages such as C++, Java, C#?

A The difference between static and non-static methods (member functions) is that static methods may access the static attributes (member variables) of a class.
B If class B inherits from class A (via public inheritance), every instance of A is considered to be an instance of B.
C If class B inherits from class A (via public inheritance), the instances of B can access all attributes (member variables) of A.
D The notions "class" and "object" mean the same thing.
*E If late binding (virtual method calls) is used, the actual method to be called is decided at run time.

Computer Networks

16 Medium Access Control (MAC) is responsible for

A the conversion between digital and analog signals in local area networks and elimination of collisions between analog and digital signals.
B redundant data transmission for error detection, and ensuring transmission repetition.
C elimination of cycles and loops in local area networks.
*D the coordination of devices accessing shared transmission media, and elimination of collisions due to concurrent transmissions.
E redundant data transmission for forward-error correction.

17 Uniform Resource Locator (URL)

A organizes domain names into dot-separated parts of maximum size of 16.

B is a description language analogous to HTML that formats any binary resource uniformly for Web browsers.

C is an application protocol used on World Wide Web (WWW).

*D must contain at least scheme, host and path. The path can be an empty string.

E is http or https followed by host, port and optional parameters.

---

18 Network Protocol

A defines physical connectors and their wiring necessary to use the network.

*B is a set of rules that defines the format and the order of messages exchanged among communicating entities, as well as the actions performed during sending/receiving such messages.

C defines only semantics, i.e. the meaning of each bit of the message exchanged among communicating entities (how particular bits should be interpreted) and nothing else.

D defines only timing, i.e. when a message can be sent and how fast the messages can be sent and nothing else.

E defines only syntax, i.e. a structure/format of data (the order in which they are presented) and nothing else.

---

19 Computers on Internet communicate via

A connection-oriented networks, where information about the state of the connection (circuit) is maintained by the network.

*B packet-switched networks, where small pieces of data (packets) are sent over the network. The packets can be routed, combined and fragmented, and no connection state is maintained by the network.

C circuit-switched networks, where the circuit may either be a fixed one or created on an "as-needed" basis.

D connection-oriented networks (circuit-switching networks), where a connection (circuit) is set up and maintained for the duration of the communication between computers.

E connection-less networks, such as the regular telephone system.

---

20 IPv6 address is 128 bits long. The standard text representation of an address

*A is split into groups of two bytes. The groups are formatted in hexadecimal and separated by colons (:). The representation can be abbreviated. E.g. FDEC::BBFF:0:FFFF.

B is split into groups of two bytes. The groups are formatted in octal and separated by colons (:). E.g. 176754:0:0:147:251:0:0

C is formatted in base64 and then split into group that are separated by colons (:).
E.g. BICw:HQAA:AAAA:AP8A:AELB: lA==.

D is split into groups of two bytes. The groups are formatted in decimal and separated by dots (.). E.g. 8193.6670.0.0.0.65280

E is split into groups of two bytes. The groups are formatted in hexadecimal and separated by semicolons (;). The leading groups that are zeros, can be omitted. E.g. FDEC;0;FFFF.

Database Systems

---

21 Choose the correct relational model that corresponds to the following E-R diagram.
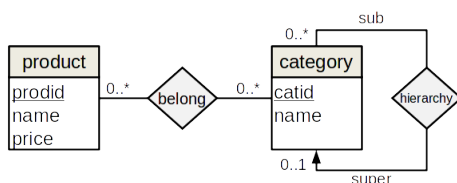
A product(prodid, name, price), category(catid, name), belong(prodid, catid), hierarchy(super_catid, sub_catid);

B product(prodid, name, price), category(catid, name);

*C product(prodid, name, price), category(catid, name, super_catid), belong(prodid, catid);

D product(prodid, name, price), category(catid, name), belong(prodid, catid), hierarchy(super_catid, sub_catid);

E product(prodid, name, price), category(catid, name), belong(prodid, catid), hierarchy(super_catid, sub_catid);

---

22 Choose the minimum required information you have to provide for *each column* when creating a new table (by a CREATE TABLE command):

A column name and integrity constraints.

B column name, primary key and integrity constraints.

C column name, data type and integrity constraints.

D column name and primary key.

*E column name and data type with size restriction when applicable.

23 Consider a relation *employee* with the schema *(id, personal_number, firstname, surname, birth_date, building, addr_street, addr_city, addr_zip, salary)* and the following functional dependencies:
id → personal_number, firstname, surname, birth_date, building, addr_street, addr_city, addr_zip, salary;
personal_number → firstname, surname, birth_date;
building → addr_street, addr_city, addr_zip;
firstname, surname, birth_date → building, salary, id.
Which one of the answers contains **all** candidate keys that are present in the relation *employee*?

A   {id}, {personal_number}
B   {personal_number}
C   {id}, {personal_number}, {building}, {firstname, surname, birth_date}
*D {id}, {personal_number}, {firstname, surname, birth_date}
E   {id}

---

24 Assume a relational database with the transaction isolation level set to serializable, the relation *employee(id, name, salary)*, and the following sequence of four SQL commands:
UPDATE employee SET salary = 0 WHERE salary < 0;
UPDATE employee SET salary = salary + 200;
SELECT name FROM employee WHERE salary > 400;
COMMIT;

This sequence is run twice in two independent transactions T1 and T2 (all commands of T1 are executed before T2 and no additional commands are executed). Choose the correct statement:

If this sequence is run twice in two independent transactions T1 and T2 (all commands of T1 are executed before T2 and no additional commands are executed), the result will be:

*A T1 returned names of all employees that had initial salary greater than 200. T2 returned names of employees that had initial salary greater than 0.
B   Both the transactions T1 and T2 returned the same results that contains names of all employees that had initial salary greater than 200.
C   Both the transactions T1 and T2 returned the same results that contains names of all employees that had initial salary greater than 0.
D   A deadlock between the transactions occurred.
E   T1 returned names of all employees that had initial salary greater than 200. T2 is blocked.

---

25 Assume the relation *employee(id, name, manager_id)* with the *manager_id* as a foreign key that references *employee.id*. There are tuples in this relation where *manager_id* **is not null**.
What does the following SQL query return?
SELECT a.name AS a, b.name AS b FROM employee a INNER JOIN employee b ON (a.manager_id = b.id)

A   Names of **all** employees in the first column, and the name of their direct manager or null value if there is not any manager in the second column.
*B Names of employees (in the first column) that have a manager whose name is displayed in the second column.
C   All possible pairs of employee names, i.e. the Cartesian (cross) product of the *employee.name* times *employee.name*.
D   Names of employees that are managers in the first column, and names of their direct subordinates in the second column (i.e. the name of the manager will repeat for each employee that he/she leads).
E   Names of **all** employees in the first column, and the name of their direct subordinates in the second column (i.e. the name of the manager will repeat for each employee that he/she leads or there will be only one tuple with null in the second column).

Software Engineering

---

26 Which one of the following lists contains types of relationships used in the UML Use Case diagram?

*A Association, Actor generalization, Include, Extend, Use case generalization
B   Association, Actor generalization, Incept, Extend, Role generalization
C   Association, Actor generalization, Include, Extend, Role generalization
D   Association, Actor generalization, Include, Extend, System generalization
E   Association, Actor generalization, Incept, Exclude, System generalization

---

27 Consider the following software development techniques:
I. Parallelization of computations
II. Load balancing among available computational nodes
III. Caching results of repeated computations

If we combine these techniques, which one of the following non-functional quality attributes will be most strongly positively influenced by them?

A Reliability
B Maintainability
C Security
D Testability
*E Performance

28 Which one is the key and statistically the most common reason why software development projects fail (to meet their budget, scope and time to market)?

A Poor knowledge of the used software development method
B Missing software design patterns
C Absence of automated testing
D Unavailability of qualified personnel
*E Poor requirement specification

29 Which one of the following descriptions best corresponds to the UML sequence diagram in the following figure?
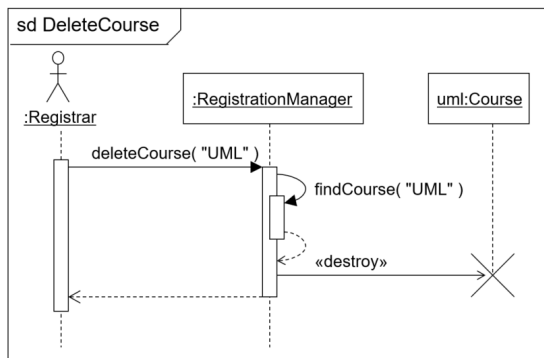
A Within the deleteCourse() method execution, an instance of the RegistrationManager class calls the findCourse() method on itself, whithin which (inside the findCourse() execution) the UML object of the Course type is deleted.
*B Within the deleteCourse() method execution, an instance of the RegistrationManager class calls the findCourse() method on itself, and then triggers the deletion of the UML object of the Course type.
C Within the deleteCourse() method execution, the RegistrationManager class calls the findCourse() method on itself, which triggers the deletion of the UML object of the Course type.
D Within the modelled scenario, three methods are executed in sequence, namely: deleteCourse(), findCourse(), destroy().
E Within the modelled scenario, two methods are executed in sequence, namely: deleteCourse(), findCourse(). That by itself invalidates the UML object of the Course class, which is deleted in effect of that.

30 Which one of the following terms does **not** have (or has the weakest) relation to agile software development?

A SCRUM
B Lean software development
C Test driven development
D Extreme programming
*E Yourdon analysis