

FI - Entrance exam - Computer Science

Jméno a příjmení - pište do okénka	Číslo přihlášky	Číslo zadání
		80

The computer science test consists of 30 questions, where you choose one of the possible answers A, B, C, D, or E. Just one answer is correct. Each correctly answered question is valued by one point, an incorrectly answered question is valued -0.25. You get zero points for multiple selected answers or no answer.

Algorithms and data structures

- 1** Consider a binary search tree without rebalancing and with unique keys. The tree is initially empty. We insert the following keys in the given order: 2, 8, 12, 1, 5, 7, 9 and then remove the key 5. What is the **preorder** of the resulting tree?
- A [8, 2, 1, 7, 9, 12]
B [8, 2, 1, 7, 12, 9]
*C [2, 1, 8, 7, 12, 9]
D [1, 2, 7, 8, 9, 12]
E [1, 7, 9, 12, 8, 2]
-
- 2** Assume an algorithm whose complexity of computation is expressed as the following function f and n is the size of input of the algorithm:
 $f(n) = n^2 + 3n + 7$ for $n < 1000$;
 $f(n) = 2f(n \operatorname{div} 2) + 6n$ for $n \geq 1000$.
(Here, $\operatorname{div} 2$ is the integer division, i.e., $\operatorname{floor}(n/2)$.)
Which of the following classes is the smallest one that contains f ?
- A $\mathcal{O}(n^2)$
B $\mathcal{O}(n)$
C $\mathcal{O}(1003007)$
D $\mathcal{O}(n^2 \log n)$
*E $\mathcal{O}(n \log n)$
-
- 3** Which statement is true?
- A We say that an algorithm is partially correct if it produces the correct output for more than half of the inputs.
*B Red-black trees are a special case of binary search trees.
C The worst-case time complexity of deleting an element from a hash table with linear probing is in $\mathcal{O}(1)$.
D The complexity of the depth-first search (DFS) algorithm on trees is linear, while on arbitrary graphs the complexity is exponential in the worst case.
E The worst-case time complexity of inserting an element into a hash table with linear probing is in $\mathcal{O}(1)$.
-
- 4** Which statement about shortest-path algorithms is true?
- A None of the other statements is true.
B If the given graph contains no negative cycles, Bellman and Ford's algorithm always terminates and returns the correct shortest paths between all pairs of vertices.
C The depth-first search (DFS) algorithm can be seen as a shortest-path algorithm, if all the edges have the same positive weight.
*D If the given graph contains no negative edges, Dijkstra's algorithm always terminates and returns the correct shortest paths from a single source vertex to all the other vertices.
E Dijkstra's algorithm can be implemented either with a queue or a stack. The choice does not change the asymptotic time complexity of the algorithm, but the queue implementation is usually faster.
-
- 5** Which statement is true for minimum binary heap?
- A The median key is stored in the root.
*B The maximum key is stored in one of the leaves.
C The maximum key is stored in the root.
D Every node either contains exactly two children or contains no children (i.e., is a leaf).
E The median key is stored in one of the leaves.

Programming

6 Consider the following function with the placeholders XXX and YYY:

```
function fun(unsigned integer n)
begin
    result = 0
    while n != 0
        digit = XXX
        result = result + digit
        n = YYY
    end while
    return result
end
```

We want the function to compute the sum of the values of all the digits in the hexadecimal representation of n . What should we write instead of XXX and YYY?

(Assume that `div` is the integer division operator and `mod` is the modulo operator.)

- A XXX = digit + 1; YYY = n - 16
- B None of the other options is correct.
- C XXX = n mod 10; YYY = n div 10
- *D XXX = n mod 16; YYY = n div 16
- E XXX = 1; YYY = n - 16

7 On common platforms, every running program has two kinds of memory available, called the **stack** and the **heap**. Which of the following three statements I, II, and III are generally true? Choose the option that contains all the true statements and none of the false ones.

- I. The stack is used to implement function calls.
- II. The heap is implemented as a binary tree.
- III. The stack is implemented as a linked list.

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

8 Consider the following code, where the print instruction outputs the given number without an end-of-line character.

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

program main()
begin
    foo(4)
end
```

What is going to be printed by the program?

- A 123124
- B The program will run forever and never halt.
- C 4321
- D 1234
- *E 1213124

9 Which statement is **false**?

- A In purely functional languages, functions can have no side effects.
- B A primitive recursive function can always be rewritten in an iterative manner.
- C A recursive function can always be rewritten in an iterative manner.
- D A tail-recursive function can always be rewritten in an iterative manner.
- *E When using call-by-value, the change of the value of a parameter inside a function can be observed outside the function.

- 10** Which statement is generally true in common OOP languages such as C++, Java, and C#?
- A The difference between a class and an object is that objects can inherit from classes but not the other way around.
 - B If a method of a class is declared as static, it means that it can access private attributes of that class; non-static methods cannot access those.
 - *C If class A inherits from class B (via public inheritance), every instance of A is considered to be an instance of B.
 - D When the early binding (non-virtual method calls) is used, the actual method to be called is decided at run time.
 - E If two or more classes inherit from the same base class, the runtime gets considerably slower, because all method calls have to check all the child classes for an implementation.

Computer Networks

- 11** Which statement is **not** valid for the link layer?
- A It is necessary to eliminate loops on the link layer, otherwise network congestion by layer-2 broadcasts is possible.
 - B Special reserved MAC addresses exist.
 - C Time-division multiplexing is not required for full-duplex communication.
 - D Using hardware from several vendors may lead to the loss of compatibility of spanning-tree protocols (STP) and, thus, form a layer-2 loop even in a properly configured network.
 - *E An ethernet frame may contain extension headers with a predetermined egress (exiting) port of the next switch.

- 12** Domain Name Space of the Internet
- A has a structure of distributed tree for routing (PRR-tree) with the maximum number of levels equal to 256.
 - *B has an inverted tree structure with a maximum number of levels equal to 128.
 - C has a virtual balanced B+-tree structure with a maximum number of levels equal to 64.
 - D has an inverted binary tree structure with the maximum number of levels equal to 1024.
 - E has a hypercube structure with a maximum number of dimensions equal to 16.

- 13** Which option forms a valid IPv4 address and network mask of a computer directly accessible from the Internet?
- A 10.0.1.127/25
 - B 192.168.1.16/28
 - C 127.0.255.255/16
 - D 224.0.0.1/24
 - *E 172.16.1.47/27

- 14** Which option provides an example of a network protocol that includes a handshake, e.g., the three-way handshake, and describes why the handshake is important and what kind of handshake is responsible for within the provided protocol?
- A Domain Name Service (DNS). The handshake assigns a unique ID to each query, thus allowing simultaneous queries.
 - B TCP and UDP. The receiver sends a negative response during the handshake to terminate the connection. Otherwise, the sender starts sending data with the congestion window set to the maximum value.
 - C UDP but only in combination with RTP and RTCP for the handshake. It sets the initial size of the congestion window to prevent congestion and allows real-time communication.
 - D Any routing protocol if routing tables are transferred by UDP. The security of the transfer is negotiated during the handshake.
 - *E A complex transport protocol, e.g., TCP. Connection parameters are negotiated during the handshake. None or negative response during the handshake means the receiver is not able to receive any data.

- 15** Which statement about the physical network layer is valid?
- A Optical cables allow higher transmission speeds. However, in comparison to metallic networks, they are also more susceptible to cross-talk among the cables.
 - *B Full duplex mode of 10BaseT standard uses one pair of conductors for transmission and one pair of conductors for simultaneous and independent receive.
 - C UTP cable consists of 8 conductors: +5V, GND, RX0, TX0, RX1, TX1, RX2, TX2.
 - D Technology known as Power Over Ethernet (POE) is not possible to operate on a gigabit network.
 - E Abbreviation UTP (network cable) stands for Unencrypted Triple Pair.

Computer systems

16 The CPU instruction SUB is used to:

- A Add two numbers.
- *B Subtract a number from another.
- C Multiply two numbers.
- D Allow input from a device.
- E Store a value in a given location.

17 A primary memory segment in a von Neumann model computer is shown in the table below. The program counter contains the data 10011.

Address	Content
10001	11001101
10010	11110001
10011	01000101
10100	11101110
10101	10010011
10110	11011010

What will be the value of the Memory Address Register (MAR)?

- *A 10011
- B 10110
- C 10001
- D 10010
- E 10101

18 One of the key features of computer architecture is the address bus. Select the proper definition of the use of such a bus.

- A This bus is used to multiplex signals among the address, control, and data buses.
- B It is a bi-directional bus used to carry signals relating to memory addresses between the processor, memory, and I/O devices.
- *C It is a uni-directional bus that carries signals relating to memory addresses between the processor and memory.
- D This bus is used to carry signals used to coordinate the computer's activities.
- E It is a bi-directional bus used to exchange data between an external device and the processor.

19 Consider three processes and their burst times in seconds: P1 with the burst time 4, P2 with the burst time 3, and P3 with the burst time 5. The time slice is 2, so every process gets executed for 2 seconds. The processes are started in the order P1, P2 and P3. What will be the cumulative waiting time of the process P2 when the Round-robin scheduling algorithm is used?

- A 2s
- B 8s
- C 7s
- D 4s
- *E 6s

20 Which option represents the integer number -13 in two's complement binary code in five bits?

- A 11011B
- B 00013B
- *C 10011B
- D 01100B
- E 11111B

Database systems

21 Select the correct statement about a foreign key (FK):

- A Only a single attribute of a given database relation can be used to define a foreign key.
- B Foreign key is an integrity constraint that evaluates a user-specified expression (only the values from the single evaluated row can be checked).
- C Foreign key is an integrity constraint that specifies that each value must be both unique and not null.
- D Foreign key cannot be a super-key of the respective database relation.
- *E A given database relation can have multiple foreign keys defined.

- 22** Consider the relation $employee(id, personal_number, name, surname, birth_date, building, address_street, address_city, address_zip, salary)$ with the following functional dependencies:
 $id \rightarrow personal_number, name, surname, birth_date, building, address_street, address_city, address_zip, salary$;
 $personal_number \rightarrow name, surname, birth_date$;
 $building \rightarrow address_street, address_city, address_zip$;
 $name, surname, birth_date \rightarrow building, salary, id$.
 Which option lists all candidate keys of the relation $employee$?
- A {personal_number}
 B {id}, {personal_number}
 C {id}, {personal_number}, {building}, {name, surname, birth_date}
 D {id}
 *E {id}, {personal_number}, {name, surname, birth_date}

- 23** Let us have a database relation $employee(id, name, surname, start_date, end_date, salary, id_supervisor)$ that represents a simplified register of the company's employees, where active employees have NULL in end_date . What is the result of the following SQL expression?

```
SELECT s.name, s.surname, SUM(e.salary)
FROM employee AS e, employee AS s
WHERE s.id = e.id_supervisor AND e.end_date IS NOT NULL
GROUP BY s.id, s.name, s.surname
```
- A The names and surnames of each employee and the total sum of salaries of all their subordinate employees that no longer work in the company.
 B The names and surnames of subordinate employees and the total sum of salaries of their respective immediate supervisors that no longer work in the company.
 C The names and surnames of employees and the total sum of salaries of their respective immediate supervisors.
 D The names and surnames of each employee and the total sum of salaries of all their subordinate employees that still work in the company.
 *E The names and surnames of employees that are supervisors and the total sum of salaries of their respective immediate subordinate employees that no longer work in the company.

- 24** Consider the following relational algebra expression:
 $documents \leftarrow \sigma_{created > '2022-01-01'}(documents) \cup \sigma_{fileSize > 0}(documents)$
 Choose the equivalent SQL expression:
- A DELETE FROM documents WHERE documents IN (SELECT * FROM documents WHERE created > '2022-01-01' UNION SELECT * FROM documents WHERE fileSize > 0);
 B UPDATE documents SET created = '2022-01-01' WHERE fileSize > 0;
 C INSERT INTO documents SELECT * FROM documents WHERE created > '2022-01-01' UNION SELECT * FROM documents WHERE fileSize > 0;
 *D DELETE FROM documents WHERE created <= '2022-01-01' AND fileSize <= 0;
 E INSERT INTO documents SELECT * FROM documents WHERE created > '2022-01-01' OR fileSize > 0;

- 25** Select the correct statement in the context of the entity-relationship (E-R) diagrams and the database relational model:
- A An entity in an E-R diagram is always represented by at least two relations in the database.
 *B Each many-to-many relationship in an E-R diagram will always be represented by one relation in the database.
 C Each one-to-many relationship in an E-R diagram will always be represented by one relation in the database.
 D There is always only one way to transform a given E-R diagram to database relations.
 E The transformation of an E-R diagram to database relations is a one-way process, so there is no way to recreate an E-R diagram from the database relations.

Software engineering

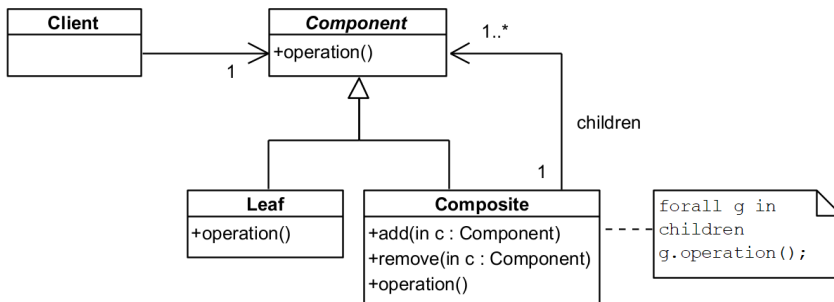
- 26** Which characteristic is **not** typical for agile development (is not part of agile manifesto)?
- A Responding to change over following a plan.
 *B Concentration on the agreed milestones over incremental delivery.
 C Individuals and interactions over processes and tools.
 D Customer collaboration over contract negotiation.
 E Working software over comprehensive documentation.

- 27** Consider the following techniques applicable to software development:
- I. The development process is organized so that faults in the system are detected and repaired before delivery to the customer.
 - II. Run-time techniques are in place to detect faults and failures, such as acceptance tests, ping/echo, and heartbeat.
 - III. The system is designed so that faults in the delivered software do not result in system failure.

Which non-functional attribute will be most positively impacted by the combination of these techniques?

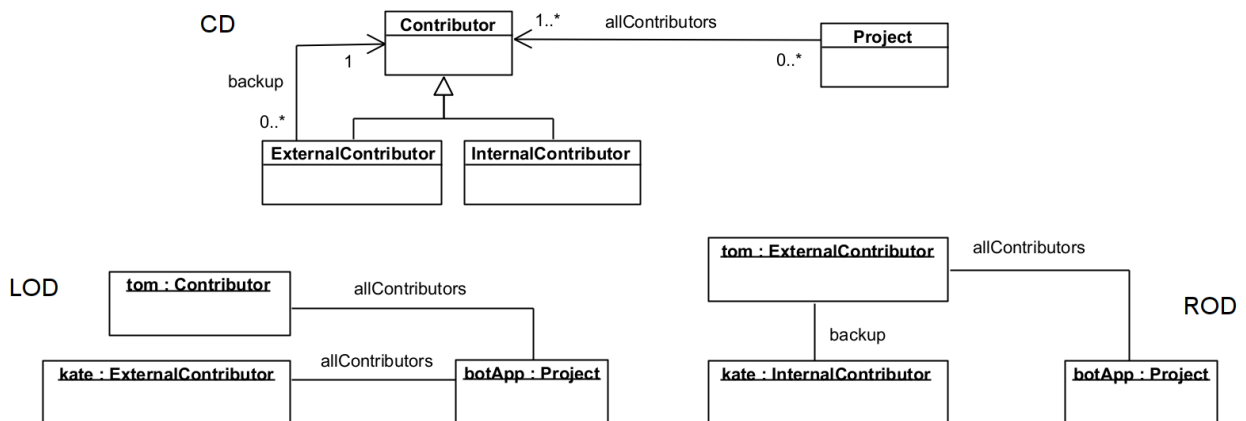
- A Security
- B Performance
- C Maintainability
- D Testability
- *E Reliability

- 28** Which statement is **not** true about the Composite design pattern depicted in the UML Class Diagram (the abstract class is in italics)?



- A A part-whole hierarchy is represented as a tree structure.
- B Both part and whole support the same operation, but while part implements it directly, the whole forwards the computation to their children.
- *C A part-whole hierarchy is represented as a network structure (not necessarily a single root).
- D A part-whole hierarchy is represented so that clients can treat the part and whole objects uniformly.
- E Composite is one of the GoF (Gangs of Four) design patterns that describe how to solve recurring design problems to design flexible and reusable object-oriented software.

- 29** Consider three UML models in the figure. Which statement about the models (i.e., the Class Diagram [CD], left Object Diagram [LOD], right Object Diagram [ROD]) is true?



- A LOD is an instance of CD. ROD is not an instance of CD because the backup connection needs to be between ExternalContributor and Contributor, which is not now.
- *B ROD is an instance of CD. LOD is not an instance of CD because there is no backup for kate.
- C ROD is an instance of CD. LOD is not an instance of CD because Project can only be connected to Contributors directly, which is not now.
- D LOD is an instance of CD. ROD is not an instance of CD because Project can only be connected to Contributors directly, which is not now.
- E None of the object diagrams (neither LOD nor ROD) is an instance of CD.

30 Consider the following statements about cloud computing:

I. Cloud offers on-demand availability of resources, especially data storage and computing power.

II. A key benefit of the cloud environment is the lower need for direct active management compared to the on-premise environment.

III. We distinguish three main service models, the Infrastructure as a service (IaaS), Platform as a service (PaaS), and Software as a service (SaaS).

Which option holds?

A I and II are true. III is false.

B I and III are true. II is false.

***C** I, II, and III are true.

D II is true. I and III are false.

E III is true. I and II are false.
