

Number: _____ Name: _____

Wrong answer to T/F questions will lead to a penalty of up to the value of the corresponding correct answer. In multiple choice questions, the penalty will be $1/(n - 1)$, with n the number of options. Penalties will be done only in the context of each question, with the minimum of 0 for the question. In each question, the first wrong answer will have no penalty.

Questão 1

For each question, answer if it is [T]rue or [F]alse or select the appropriate answer, in multiple choice questions.

1. F In a distributed system, when a component does not reply it is possible to determine if the fault has occurred in the component or in the communication system. (T/F)
2. T It is impossible to a process in a distributed system to know the global state of the system at some given time instant. (T/F)
3. F A distributed system is totally transparent when it exposes to its users the number, location and nature of the components. (T/F)
4. A The main advantage of geo-replication is: (A) to reduce the latency experienced by clients; (B) improve the fault tolerance; (C) contribute to distribute the load; (D) none of the previous ones. (A/B/C/D)
5. T Running a component of a distributed system in a *container* is more efficient than running it in a virtual machinel. (T/F)
6. T To mask faults in a distributed system it is common to rely on some form of redundancy. (T/F)
7. A In an architecture of 3 layers, what is the level that is typically easier to replicate: (A) the application layer; (B) the storage layer; (C) both have the same complexity. (A/B/C)
8. T In a distributed system, placing the *proxy* close to the client or to the server leads to distinct properties for the system. (T/F)
9. T A aptitioned server, compared to a replicated server, tends to receive a lower number of writes. (T/F)
10. C When combining partitioning and replication, with n machines and n partitions, and each partition replicated in k machines, for r read operations and w write operations, each machine processes on average the given number of operations: (A) $r / n + w / n$; (B) $(r + w) / n$; (C) $r / n + w.k / n$; (D) $(r + w) . k / n$; (E) none of the previous. (A/B/C/D/E)
11. F In a structured P2P system, the physical address of nodes (e.g., IP) can have no impact in the location of each node in the topology of the network. (T/F)
12. F In Bittorrent system, nodes that are transferring the same file establish a structured P2P network. (T/F)

Questão 2

For each question, answer if it is [T]rue or [F]alse or select the appropriate answer, in multiple choice questions.

13. T Web Sockets allow to notify Web client of the occurrence of an event - e.g., in an email application, Web Sockets can be used to notify the existence of a new message. (T/F)
14. F When a process sends an UDP message with 100 bytes it is possible that this message is received in two successive *receive* operation, with each one receiving 50 bytes of the message. (T/F)
15. F TCP is a volatile communication protocol. (T/F)
16. T In the support for SOAP available in Java (JAX-WS, Jakarta), WSDL is automatically generated by the server, based on the annotations in the server's methods (or in the server's interface). (T/F)
17. T Serializing messages using the format of the sender (*receiver makes it right*) is more efficient than using an independent intermediate format. (T/F)
18. T The serialization mechanism of Java, when serializing object O , the serialized state includes the name of the class of object O . (T/F)
19. C What is the main advantage of ProtoBuf serialization when compared to JSON: (A) small messages; (B) faster to serialize/deserialize; (C) both of the previous two; (D) none of the previous. (A/B/C/D)
20. T In a remote invocation mechanism, an object is typically passed by value. (T/F)

21. F Even when experiencing temporary network faults, the distributed *garbage-collection* mechanism of Java guarantees that a remote object (server) is never garbage-collected while there is a reference to the remote object. (T/F)
22. T In Java, if a server does not invoke another remote server when processing a client request, a correct (even if not very efficient) way to do concurrency control is to define all methods as *synchronized*. (V/F)
23. F In REST, it is considered good practice to return code 400 BAD REQUEST for an read, update or delete operation that references a non-existing resource. (T/F)
24. T The fact that each resource has an URL allows to use the existing caching mechanisms for HTTP protocol. (T/F)

Questão 3

To implement a remote method invocation protocol with semantics *at most once* or *exactly once*, it is possible to use an algorithm with retransmissions (in which the client send the request more than once) and that filters duplicates (in which the server checks if it has already received the request).

- a When operations are idempotent, it is necessary to store the result of the invocation in the server? Justify.

No. As the result of executing an idempotent operation more than once is the same as executing it only once, if the client sends the operation again, the server can simply execute the method again and return the result.

- b Using sequence numbers to identify operation allows the server an optimization that consists in keeping only the result of the last operation for each client. If clients use instead of sequence number, the time of the clock (in nanoseconds), would it be possible for the server to perform the same optimization? Justify explaining how or why it is not possible.

Yes, because... / No, because...

Yes, assuming that a client cannot generate two operations at the same exact nanosecond. The identifier of the operation would be the timestamp and the identifier of the client (as it is also necessary when using sequence numbers). If using a stream channel (as TCP), this is enough. If using messages (as UDP), where messages can be reorder during transmission, an operation would need to include the timestamp of the operation and the timestamp of the previous operation, so that the server could know if it has missed some operation.

[This is the full answer, but a reply considering a TCP channel only would be ok.]

Questão 4

The restrictions associated to the global pandemics led clients of many entities (physical services, such as banks, hairdressers, etc.) to need to make an appointment before going to the entity. Consider you want to create a system to manage the reservations of entities, that will allow any entity (physical service) to make available, for each day, a set of periods for appointments that can be reserved by their clients. The system will support entities located at any place in the world.

Some of the operations available in this system are the following:

- (1) create an entity;
 - (2) create a new appointment period (e.g., specify that it is possible to schedule an appointment at 9h00);
 - (3) reserve an appointment period previously created (e.g., client reserves the appointment period of 9h00);
 - (4) cancel a reservation previously done (e.g., client cancels the reservation of appointment period at 9h00, making it available again);
 - (5) get the list of appointment periods for an entity, in a given date;
 - (6) remove an appointment period;
 - (7) remove an entity.
- a) Consider that you will use the *cloud* to implement the system, and you have access to five data centers - Europe, Africa, North America, South America and Asia. Discuss how you would distribute the data of the system by the data center – which data would be stored in each data center. (Avoid generic sentences - reply for the concrete example of the question.)

As appointments are for physical services, one can expect users and services to be in the same geographical area - it should not be very common for a portuguese person to schedule a haircut in New York. (note that if the service takes a little longer in some rare occasion, that is perfectly ok)

As such, data for entities (including appointments) should be located in the closest data center. For tolerating fault in case a disaster occurs that makes the data center unavailable, this data should be replicated in other data center.

- b) Would it make sense to use a CDN system to complete the solution proposed in the previous question, assuming that the CDN system would have a point-of-presence in each Internet provider (e.g., in Portugal, it would be Meo, Vodafone, Nos, etc.)? Justify discussing which operation could benefit from using the CDN and which ones cannot be supported by the CDN. (Avoid generic sentences - reply for the concrete example of the question.)

A CDN would not add much to this system, as data is changing very frequently and one can expect that it will not be accessed very frequently.

Note that CDNs can also be used for supporting mutable data, but the fact that there are several PoP in each area requires lots of synchronization for keeping the data in the PoPs synchronized. If we consider that the number of accesses will be low, it is hardly justifiable the complexity.

- c) Given that you intend to use domain “reservationsforall.com” for your system, and that you want to make the operations available using REST, specify which URL and operation would you use for operation “(3) reserve an appointment period previously created”. Explain why you propose such URL, discussing the resources of the system and possible alternatives.

`http://reservationsforall/rest/entities/{entityid}/{appointment_period}`

In REST, an URL should identify a resource. We can consider two main resources in this application. Entities, that represent the entity providing the service.

Appointment periods, that represent an appointment period made available by an entity. As such, it makes sense that we consider the entity the topmost resource in the URL for operation (3) - an entity can exist without appointments, but an appointment cannot exist without an entity.

The third resource in the system would be the user, but the combination `{entityid}/{appointment_period}/{userid}` is not a resource, so it should not be used in the URL.