Roll No: _____

Amrita Vishwa Vidyapeetham

M.Tech Second Assessment – February 2013
Second Semester

Embedded Systems

ES623 Networked Embedded Systems Answer Key

Time: Two Hours **Maximum: 50 Marks**

Answer all Questions

---

1. Discuss the consequences of an error in the external clock synchronization. What effect can such an error have on the internal clock synchronization in the worst possible scenario? **(4 marks)**

Solution:

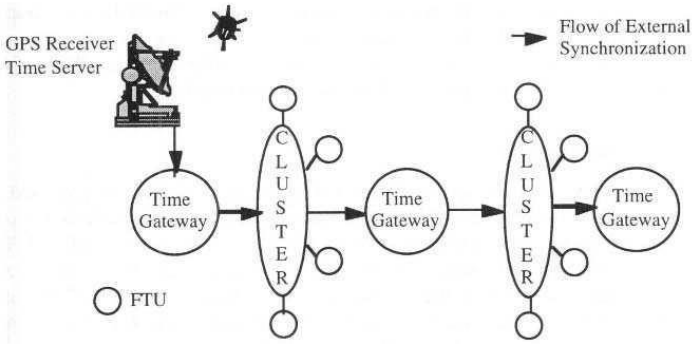

Time gateway must synchronize the global time of its cluster with the time received from the time server. This synchronization is unidirectional, and therefore asymmetric.

Internal synchronization is a cooperative activity among all the members of a cluster whereas external synchronization is an authoritarian process: the time server forces its view of external time on all its subordinates.

A time gateway will only accept an external synchronization message if its content is sufficiently close to its view of the external time. The time server has only a limited authority to correct the drift rate of a cluster.

The implementation must guarantee that it is impossible for a faulty external synchronization to interfere with the proper operation of the internal synchronization, i.e., with the generation of global time within a cluster.

The worst possible failure scenario occurs if the external time server fails maliciously. This leads to a common-mode deviation of the global time from the external time base with the maximum permitted correction rate.

The internal synchronization within a cluster will not be affected by this controlled drift from the external time base.

2. Consider that there is a difference in *endianness*, i.e., the byte ordering of data, between the sender and the receiver of a message. The sender assumes *big endian*, i.e., the most significant byte is first, and the receiver assumes *little endian*, i.e., the least significant byte is first. How is this property mismatch resolved in a computer network? **(3 marks)**

Solution:

Property mismatch is resolved by a resource controller, i.e., the gateway component. The resource controller transforms the information from the representation used in one subsystem to that used in the other subsystem.

3. What is thrashing? What mechanisms can lead to thrashing? How should you react in an event-triggered system if thrashing is observed? **(5 marks)**

Solution:

Throughput of a system decreasing abruptly with increasing load, is called *thrashing.*

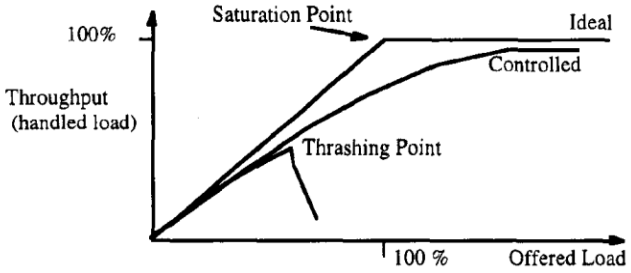

The throughput increases with increasing load until the saturation point has been reached. Thereon, the throughput remains constant. A system has a *controlled* load-throughput characteristic if the throughput increases monotonically with the load and reaches the maximum throughput asymptotically. If the throughput increases up to a certain point, the *thrashing point*, and thereafter decreases abruptly, then, we say the system is *thrashing.*

 # Mechanisms that can Cause Thrashing:

(i) The retry mechanism in the PAR protocol: If a communication system slows down because it can barely handle the offered load, a high-level PAR protocol reaches its time-outs, and generates *additional* load.

(ii) Operating system services: In a dynamic scheduling environment, the time needed to find a feasible schedule increases more than linearly as the offered load reaches the capacity limit. This increase in the amount of scheduling overhead further decreases the computational resources that are available for the application tasks.

The *only* successful technique to avoid thrashing in explicit flow-control schemes is to monitor the resource requirements of the system continuously and to exercise a stringent back-pressure flow control as soon as a decrease in the throughput is observed.

4. Consider a bus system without global time where a token protocol controls media access to the bus. The token protocol has a maximum token rotation time *TRT* of *10* msec. Calculate the latency jitter of a high level PAR protocol that allows three retries, assuming that the lower level protocol used for this implementation has a *dmin* of *1* msec and a *dmax* of *30* msec. Calculate the error detection latency at the sender if the time out is set to at least *12* msec. **(3 marks)**

Solution:

PAR jitter = 30 – 1 = 29msec

Error detection latency = 3 x 12mSec = 36msec

5. *The JAS 39 Gripen is a fighter aircraft manufactured by the Swedish aerospace company Saab. On 2 February 1989, the first prototype JAS 39-1 crashed on its sixth flight, when attempting to land in Linkoping. The cause of the crash was identified as pilot induced oscillation which was due to problems with the flight control system. Extremely gusty winds were also a contributing factor. The software was unable to handle strong winds at low speeds, whereas the plane itself responded too slowly to the pilot's controls.*

   Identify the flow control involved in the above real-time system. How fault tolerance is implemented in such systems? Discuss the process interface between the controlled object and the computer system, in view of flow control. **(5 marks)**

   Solution:

   # Explicit flow control is involved in the above real-time systems.

   # In *explicit flow control,* the receiver sends an explicit acknowledgment message to the sender, informing the sender that the sender's previous message arrived correctly, and that the receiver is now ready to accept the next message.

   # Positive-Acknowledgment-or-Retransmission (PAR) protocol.

   # From the point of view of flow control, the most critical interface in a real-time system is the process interface between the controlled object and the computer system.

   # It cannot be assumed that *all* events occurring in the controlled object are in the sphere of control of the computer system. If, in an event-triggered system, more events occur in the controlled objects, then, the computer system may be overloaded by such an "event shower", and thereby miss important deadlines.

   # When the subsystem uses implicit flow control, the consumer subsystem with explicit flow control can consume information only at the speed determined by its receivers.

   # In order not to lose any information, adequate buffering must be provided at this interface. Determining the proper buffer size is a delicate design issue that is often ignored at the design level, and left to the programmer at the end of the line.

6. Fault tolerance can be implemented by two, fail-silent components or by TMR. Discuss the advantages and disadvantages of each one of these methods. **(5 marks)**

   Solution:

   # Fail-Silent Nodes: A fail-silent node must detect all internal failures within a short latency, and must map these failures to a single external failure mode, a fail-silent node failure.

   The advantages of FTU with shadow nodes are:

   (i) Whenever an active node fails, the redundancy within the FTU is reestablished within a short time interval.

   (ii) During normal operation the shadow node does not consume any bandwidth of the communication system.

   (iii) During repair of the failed node, the redundancy within the FTU is maintained.

In *exact voting*, a bit-by-bit comparison of the data fields in the redundant result messages is performed. Exact voting requires replica determinate computational channels.
In *inexact voting* two messages are assumed to contain the *same* result if the results are within some application-specific interval. Inexact voting must be used if replica determinism of the replicated nodes cannot be guaranteed.

7. What is the difference between a state observation and an event observation? Discuss their advantages and disadvantages. **(4 marks)**

Solution:

An observation is a *state observation* if the value of the observation contains the absolute state of the RT entity. The time of the state observation refers to the point in real time when the RT entity was sampled.
An observation is an *event observation* if it contains information about the *change of value* between the "old state" and the "new state". The time of the event observation denotes the best estimate of the point in time of this event.

8. Describe the elements of an interface. What are the characteristics of world interfaces and message interfaces? **(4 marks)**

Solution:

An interface is a common boundary between two subsystems. A correctly designed interface provides understandable abstractions, to the interfacing partners, which capture the essential properties of the interfacing subsystems and hide the irrelevant details. An interface between two subsystems of a real-time system can be characterized by:
(i) The *control properties*, i.e., the properties of the control signals crossing the interface, e.g., which task must be activated if a particular event happens.
(ii) The *temporal properties*, i.e., the temporal constraints that must be satisfied by the control signals and by the data that cross the interface.
(iii) The *functional intent*, i.e., the specification of the intended functions of the interfacing partner.
(iv) The *data properties*, i.e., the structure and semantics of the data elements crossing the interface.

Characteristic	Concrete World Interface	Abstract Message Interface
Information Representation	unique, determined by the given device	uniform within the whole cluster
Coupling	tight, determined by the specific I/O protocol of the connected device	weaker, determined by the message communication protocol
Coding	analog or digital, unique	digital, uniform codes
Time-base	dense	possibly sparse
Inaterconnetion Pattern	one-to-one	one-to-many
Freedom in Design	determined by the format and timing of the physical I/O devices	determined by the uniform standards of the architecture

9. What are the temporal obligations of clients and servers at a client-server interface in a real-time system? **(4 marks)**

Solution:

Three temporal parameters characterize such a clientserver interaction:

(i) The maximum response time, RESP, that is expected by the client, and stated in the specification,

(ii) The worst-case execution time, WCET, of the server that is determined by the implementation of the server, and

(iii) The minimum time, MINT, between two successive requests by the client. It is important to note that the WCET is in the sphere of control of the server, and that the minimum time between two successive requests, MINT, is in the sphere of control of the client. In a hard real-time environment, the implementation must guarantee that the condition

WCET < RESP

holds, under the assumption that the client respects its obligation to keep a minimum temporal distance MINT between two successive requests.

10. From the report of The President's Commission on the Accident at Three Mile Island, October 1979, Washington, D.C.:

On March 28, 1979, the United States experienced the worst accident in the history of commercial nuclear power generation. Two weeks later, the President of the United States, Mr. Jimmy Carter, established a 12-member Presidential Commission.

... At 4:00 a.m. on March 28, 1979, a serious accident occurred at the Three Mile Island 2 nuclear power plant near Middletown, Pennsylvania. The accident was initiated by mechanical malfunctions in the plant and made much worse by a combination of human errors in responding to it.

.... The pilot-operated relief valve (PORV) at the top of the pressurizer opened as expected when pressure rose but failed to close when pressure decreased, thereby creating an opening in the primary coolant system -- a small-break loss-of-coolant accident (LOCA). The PORV indicator light in the control room showed only that the signal had been sent to close the PORV rather than the fact that the PORV remained open. The operators, relying on the indicator light and believing that the PORV had closed, did not heed other indications and were unaware of the PORV failure; the LOCA continued for over 2 hours.

Suggest few requirements of a communication system for such a safety critical application. Discuss how these mechanical and human errors can be handled and managed in distributed real-time systems. **(7 marks)**

Solution:

Errors that occur during the message transmission must be detected, and should be corrected without increasing the jitter of the protocol latency. If an error cannot be corrected, all the communicating partners, the sender and all the receivers, must be informed about the occurrence of the error with a low latency.

The failure of a node must be detected by the communication protocol, and must be reported consistently to all the remaining nodes of the ensemble. In real-time systems, the prompt and consistent detection of node failures at both the receiver and at the sender is important which is the function of the membership service.

In a real-time system, the end-to-end acknowledgment about the success or failure of a communication action can arise from a node that is different from the receiver of an output message. An output message to an actuator in the environment should cause some effect in the environment which is monitored by an independent sensor.

11. The following is a portion of an article titled 'Electronic car bugs: What drivers need to know' from the *New Scientist:*
During the 1980s, drivers of Mercedes-Benz cars with anti-lock brakes (ABS) reported that their brakes were failing on a section of autobahn in the Saarland region of Germany. The problem, caused by electromagnetic interference (EMI) from a nearby radio transmitter, was solved by putting up a giant wire mesh by the side of the road to shield traffic from its radio transmissions. There have also been documented cases of EMI causing problems with remote locking and security systems.
How would you classify this fault? State reasons. **(3 marks)**

Solution:

Transient failure
If the system continues to operate after the failure, then the failure is a *transient failure*

12. What is the difference between temporal control and logical control? Give example. **(3 marks)**

Solution:

 # Temporal control is concerned with the determination of the points in time when a task must be activated or when a task must be blocked, because some conditions *outside* the task are not satisfied at a particular moment.

 # Logical control is concerned with the control flow *within* a task that is determined by the given program structure and the particular input data to achieve the desired data transformation.
