

Homework Assignment 1

Practice the following questions based on chapters 1, 2, and 3 of the textbook.

Chapter 1

Multiple Choice

1. In what way is an operating system like a government?
A) It seldom functions correctly.
B) It creates an environment within which other programs can do useful work.
C) It performs most useful functions by itself.
D) It is always concerned primarily with the individual's needs.

Ans: B

2. ____ operating systems are designed primarily to maximize resource utilization.
A) PC
B) Handheld computer
C) Mainframe
D) Network

Ans: C

3. The most common secondary storage device is ____.
A) random access memory
B) solid state disks
C) tape drives
D) magnetic disk

Ans: D

4. Which of the following would lead you to believe that a given system is an SMP-type system?
A) Each processor is assigned a specific task.
B) There is a boss-worker relationship between the processors.
C) Each processor performs all tasks within the operating system.
D) None of the above

Ans: C

5. A ____ can be used to prevent a user program from never returning control to the operating system.
A) portal
B) program counter
C) firewall

D) timer

Ans: D

6. Embedded computers typically run on a ____ operating system.

- A) real-time
- B) Windows XP
- C) network
- D) clustered

Ans: A

7. Bluetooth and 802.11 devices use wireless technology to communicate over several feet, in essence creating a ____.

- A) local-area network
- B) wide-area network
- C) small-area network
- D) metropolitan-area network

Ans: C

8. A clustered system ____.

- A) gathers together multiple CPUs to accomplish computational work
- B) is an operating system that provides file sharing across a network
- C) is used when rigid time requirements are present
- D) can only operate one application at a time

Ans: A

9. Which of the following is a property of peer-to-peer systems?

- A) Clients and servers are not distinguished from one another.
- B) Separate machines act as either the client of the server but not both.
- C) They do not offer any advantages over traditional client-server systems.
- D) They suffer from the server acting as the bottleneck in performance.

Ans: A

10. Two important design issues for cache memory are ____.

- A) speed and volatility
- B) size and replacement policy
- C) power consumption and reusability
- D) size and access privileges

Ans: B

11. What are some other terms for kernel mode?

- A) supervisor mode
- B) system mode
- C) privileged mode
- D) All of the above

Ans: D

12. Which of the following statements concerning open source operating systems is true?

- A) Solaris is open source.
- B) Source code is freely available.
- C) They are always more secure than commercial, closed systems.
- D) All open source operating systems share the same set of goals.

Ans: B

13. Which of the following operating systems is not open source?

- A) Windows
- B) BSD UNIX
- C) Linux
- D) PCLinuxOS

Ans: A

14. A _____ provides a file-system interface which allows clients to create and modify files.

- A) compute-server system
- B) file-server system
- C) wireless network
- D) network computer

Ans: B

15. A _____ is a custom build of the Linux operating system

- A) LiveCD
- B) installation
- C) distribution
- D) VMWare Player

Ans: C

16. _____ is a set of software frameworks that provide additional services to application developers.

- A) System programs
- B) Virtualization
- C) Cloud computing
- D) Middleware

Ans: D

17. What statement concerning privileged instructions is considered false?

- A) They may cause harm to the system.
- B) They can only be executed in kernel mode.
- C) They cannot be attempted from user mode.
- D) They are used to manage interrupts.

Ans: C

18. Which of the following statements is false?

- A) Mobile devices must be concerned with power consumption.
- B) Mobile devices can provide features that are unavailable on desktop or laptop computers.
- C) The difference in storage capacity between a mobile device and laptop is shrinking.
- D) Mobile devices usually have fewer processing cores than a standard desktop computer.

Ans: C

19. A(n) _____ is the unit of work in a system.

- A) process
- B) operating system
- C) timer
- D) mode bit

Ans: A

20. The two separate modes of operating in a system are

- A) supervisor mode and system mode
- B) kernel mode and privileged mode
- C) physical mode and logical mode
- D) user mode and kernel mode

Ans: D

Essay

21. Explain why an operating system can be viewed as a resource allocator.

Ans: A computer system has many resources that may be required to solve a problem: CPU time, memory space, file-storage space, I/O devices, and so on. The operating system acts as the manager of these resources. Facing numerous and possibly conflicting requests for resources, the operating system must decide how to allocate them to specific programs and users so that it can operate the computer system efficiently and fairly.

22. Explain the purpose of an interrupt vector.

Ans: The interrupt vector is merely a table of pointers to specific interrupt-handling routines. Because there are a fixed number of interrupts, this table allows for more efficient handling of the interrupts than with a general-purpose, interrupt-processing routine.

23. What is a bootstrap program, and where is it stored?

Ans: A bootstrap program is the initial program that the computer runs when it is powered up or rebooted. It initializes all aspects of the system, from CPU registers to device controllers to memory contents. Typically, it is stored in read-only memory (ROM) or electrically erasable programmable read-only memory (EEPROM), known by the general term firmware, within the computer hardware.

24. What role do device controllers and device drivers play in a computer system?

Ans: A general-purpose computer system consists of CPUs and multiple device controllers that are connected through a common bus. Each device controller is in charge of a specific type of device. The device controller is responsible for moving the data between the peripheral devices that it controls and its local buffer storage. Typically, operating systems have a device driver for each device controller. This device driver understands the device controller and presents a uniform interface for the device to the rest of the operating system.

25. Why are clustered systems considered to provide high-availability service?

Ans: Clustered systems are considered high-availability in that these types of systems have redundancies capable of taking over a specific process or task in the case of a failure. The redundancies are inherent due to the fact that clustered systems are composed of two or more individual systems coupled together.

26. Describe the differences between physical, virtual, and logical memory.

Ans: Physical memory is the memory available for machines to execute operations (i.e., cache, random access memory, etc.). Virtual memory is a method through which programs can be executed that requires space larger than that available in physical memory by using disk memory as a backing store for main memory. Logical memory is an abstraction of the computer's different types of memory that allows programmers and applications a simplified view of memory and frees them from concern over memory-storage limitations.

27. Describe the operating system's two modes of operation.

Ans: In order to ensure the proper execution of the operating system, most computer systems provide hardware support to distinguish between user mode and kernel mode. A mode bit is added to the hardware of the computer to indicate the current mode: kernel (0) or user (1). When

the computer system is executing on behalf of a user application, the system is in user mode. However, when a user application requests a service from the operating system (via a system call), it must transition from user to kernel mode to fulfill the request.

28. Explain cache coherency.

Ans: In multiprocessor environments, two copies of the same data may reside in the local cache of each CPU. Whenever one CPU alters the data, the cache of the other CPU must receive an updated version of this data. Cache coherency involves ensuring that multiple caches store the most updated version of the stored data.

29. Why is main memory not suitable for permanent program storage or backup purposes? Furthermore, what is the main disadvantage to storing information on a magnetic disk drive as opposed to main memory?

Ans: Main memory is a volatile memory in that any power loss to the system will result in erasure of the data stored within that memory. While disk drives can store more information permanently than main memory, disk drives are significantly slower.

30. Describe the compute-server and file-server types of server systems.

Ans: The compute-server system provides an interface to which a client can send a request to perform an action (for example, read data from a database); in response, the server executes the action and sends back results to the client. The file-server system provides a file-system interface where clients can create, update, read, and delete files. An example of such a system is a Web server that delivers files to clients running Web browsers.

31. Computer systems can be divided into four approximate components. What are they?

Ans: Hardware, operating system, application programs, and users.

32. Distinguish between system and application programs.

Ans: System programs are not part of the kernel, but still are associated with the operating system. Application programs are not associated with the operating of the system.

33. Describe why direct memory access (DMA) is considered an efficient mechanism for performing I/O.

Ans: DMA is efficient for moving large amounts of data between I/O devices and main memory. It is considered efficient because it removes the CPU from being responsible for transferring data. DMA instructs the device controller to move data between the devices and main memory.

34. Describe why multi-core processing is more efficient than placing each processor on its own chip.

Ans: A large reason why it is more efficient is that communication between processors on the same chip is faster than processors on separate chips.

35. Distinguish between uniform memory access (UMA) and non-uniform memory access (NUMA) systems.

Ans: On UMA systems, accessing RAM takes the same amount of time from any CPU. On NUMA systems, accessing some parts of memory may take longer than accessing other parts of

36. Explain the difference between singly, doubly, and circularly linked lists.

Ans: A singly linked list is where each item points to its successor. A doubly linked list allows an item to point to its predecessor or successor. A circularly linked list is the where the last element points back to the first.

37. What two operating systems currently dominate mobile computing?

Ans: Apple's iOS and Google's Android

38. Explain the difference between protection and security.

Ans: Protection is concerned with controlling the access of processes or users to the resources of the computer system. The role of security is to defend the system from internal or external attacks.

39. Distinguish mobile computing from traditional desktop computing.

Ans: Mobile computing takes place on handheld devices and tablets. Because these devices are portable and lightweight, they typically do not have the processing power and storage capacity of desktop systems. However, features such as GPS and accelerometers have allowed mobile devices to provide functionality that is unavailable to desktop systems.

40. Describe cloud computing.

Ans: Cloud computing is a type of computing that delivers computing, storage, and application services across a network. Cloud computing often uses virtualization to provide its functionality. There are many different types of cloud environments, as well as services offered. Cloud computing may be either public, private, or a hybrid of the two. Additionally, cloud computing may offer applications, platforms, or system infrastructures.

True/False

41. The operating system kernel consists of all system and application programs in a computer.

Ans: False

42. Flash memory is slower than DRAM but needs no power to retain its contents.

Ans: True

43. A system call is triggered by hardware.

Ans: False

44. UNIX does not allow users to escalate privileges to gain extra permissions for a restricted activity.

Ans: False

45. Processors for most mobile devices run at a slower speed than a processor in a desktop PC.

Ans: True

46. Interrupts may be triggered by either hardware or software

Ans: True

47. A dual-core system requires each core has its own cache memory.

Ans: False

48. Virtually all modern operating systems provide support for SMP

Ans: True

49. All computer systems have some sort of user interaction.

Ans: False

50. Solid state disks are generally faster than magnetic disks.

Ans: True

51. Solid state disks are considered volatile storage.

Ans: False

52. There is no universally accepted definition of an operating system.

Ans: True

Chapter 2

Multiple Choice

1. A _____ is an example of a systems program.

- A) command interpreter
- B) Web browser
- C) text formatter
- D) database system

Ans: A

2. If a program terminates abnormally, a dump of memory may be examined by a _____ to determine the cause of the problem.

- A) module
- B) debugger
- C) shell
- D) control card

Ans: B

3. A message-passing model is _____.

- A) easier to implement than a shared memory model for intercomputer communication
- B) faster than the shared memory model
- C) a network protocol, and does not apply to operating systems
- D) only useful for small simple operating systems

Ans: A

4. Policy _____.

- A) determines how to do something
- B) determines what will be done
- C) is not likely to change across places
- D) is not likely to change over time

Ans: B

5. The major difficulty in designing a layered operating system approach is ____.
- A) appropriately defining the various layers
 - B) making sure that each layer hides certain data structures, hardware, and operations from higher-level layers
 - C) debugging a particular layer
 - D) making sure each layer is easily converted to modules

Ans: A

6. A microkernel is a kernel ____.
- A) containing many components that are optimized to reduce resident memory size
 - B) that is compressed before loading in order to reduce its resident memory size
 - C) that is compiled to produce the smallest size possible when stored to disk
 - D) that is stripped of all nonessential components

Ans: D

7. To the SYSGEN program of an operating system, the least useful piece of information is ____.
- A) the CPU being used
 - B) amount of memory available
 - C) what applications to install
 - D) operating-system options such as buffer sizes or CPU scheduling algorithms

Ans: C

8. A boot block ____.
- A) typically only knows the location and length of the rest of the bootstrap program
 - B) typically is sophisticated enough to load the operating system and begin its execution
 - C) is composed of multiple disk blocks
 - D) is composed of multiple disk cylinders

Ans: A

9. ____ provide(s) an interface to the services provided by an operating system.
- A) Shared memory
 - B) System calls
 - C) Simulators
 - D) Communication

Ans: B

10. ____ is not one of the major categories of system calls.
- A) Process control
 - B) Communications

- C) Protection
- D) Security

Ans: D

11. _____ allow operating system services to be loaded dynamically.
- A) Virtual machines
 - B) Modules
 - C) File systems
 - D) Graphical user interfaces

Ans: B

12. Microkernels use _____ for communication.
- A) message passing
 - B) shared memory
 - C) system calls
 - D) virtualization

Ans: A

13. The Windows `CreateProcess()` system call creates a new process. What is the equivalent system call in UNIX:
- A) `NTCreateProcess()`
 - B) `process()`
 - C) `fork()`
 - D) `getpid()`

Ans: C

14. The `close()` system call in UNIX is used to close a file. What is the equivalent system call in Windows:
- A) `CloseHandle()`
 - B) `close()`
 - C) `CloseFile()`
 - D) `Exit()`

Ans: A

15. The Windows `CreateFile()` system call is used to create a file. What is the equivalent system call in UNIX:
- A) `ioctl()`
 - B) `open()`
 - C) `fork()`
 - D) `createfile()`

Ans: B

16. Android runs Java programs _____

- A) in the Dalvik virtual machine.
- B) natively.
- C) in the Java virtual machine.
- D) Android does not run Java programs.

Ans: A

17. _____ is a mobile operating system designed for the iPhone and iPad.

- A) Mac OS X
- B) Android
- C) UNIX
- D) iOS

Ans: D

18. The _____ provides a portion of the system call interface for UNIX and Linux.

- A) POSIX
- B) Java
- C) Standard C library
- D) Standard API

Ans: C

19. Which of the following statements is incorrect?

- A) An operating system provides an environment for the execution of programs.
- B) An operating system manages system resources.
- C) Operating systems provide both command line as well as graphical user interfaces.
- D) Operating systems must provide both protection and security.

Ans: C

20. _____ is/are not a technique for passing parameters from an application to a system call.

- A) Cache memory
- B) Registers
- C) Stack
- D) Special block in memory

Ans: A

Essay

21. There are two different ways that commands can be processed by a command interpreter. One way is to allow the command interpreter to contain the code needed to execute the command. The other way is to implement the commands through system programs. Compare and contrast the two approaches.

Ans: In the first approach, upon the user issuing a command, the interpreter jumps to the appropriate section of code, executes the command, and returns control back to the user. In the second approach, the interpreter loads the appropriate program into memory along with the appropriate arguments. The advantage of the first method is speed and overall simplicity. The disadvantage to this technique is that new commands require rewriting the interpreter program which, after a number of modifications, may get complicated, messy, or too large. The advantage to the second method is that new commands can be added without altering the command interpreter. The disadvantage is reduced speed and the clumsiness of passing parameters from the interpreter to the system program.

22. Describe the relationship between an API, the system-call interface, and the operating system.

Ans: The system-call interface of a programming language serves as a link to system calls made available by the operating system. This interface intercepts function calls in the API and invokes the necessary system call within the operating system. Thus, most of the details of the operating-system interface are hidden from the programmer by the API and are managed by the run-time support library.

23. Describe three general methods used to pass parameters to the operating system during system calls.

Ans: The simplest approach is to pass the parameters in registers. In some cases, there may be more parameters than registers. In these cases, the parameters are generally stored in a block, or table, of memory, and the address of the block is passed as a parameter in a register. Parameters can also be placed, or pushed, onto the stack by the program and popped off the stack by the operating system.

24. What are the advantages of using a higher-level language to implement an operating system?

Ans: The code can be written faster, is more compact, and is easier to understand and debug. In addition, improvements in compiler technology will improve the generated code for the entire operating system by simple recompilation. Finally, an operating system is far easier to port — to move to some other hardware — if it is written in a higher-level language.

25. Describe some requirements, or goals, when designing an operating system.

Ans: Requirements can be divided into user and system goals. Users desire a system that is convenient to use, easy to learn, and to use, reliable, safe, and fast. System goals are defined by those people who must design, create, maintain, and operate the system: The system should be easy to design, implement, and maintain; it should be flexible, reliable, error-free, and efficient.

26. What are the advantages and disadvantages of using a microkernel approach?

Ans: One benefit of the microkernel approach is ease of extending the operating system. All new services are added to user space and consequently do not require modification of the kernel. The microkernel also provides more security and reliability, since most services are running as user — rather than kernel — processes. Unfortunately, microkernels can suffer from performance decreases due to increased system function overhead.

27. Explain why a modular kernel may be the best of the current operating system design techniques.

Ans: The modular approach combines the benefits of both the layered and microkernel design techniques. In a modular design, the kernel needs only to have the capability to perform the required functions and know how to communicate between modules. However, if more functionality is required in the kernel, then the user can dynamically load modules into the kernel. The kernel can have sections with well-defined, protected interfaces, a desirable property found in layered systems. More flexibility can be achieved by allowing the modules to communicate with one another.

28. Describe how Mac OS X is considered a hybrid system.

Ans: Primarily because the kernel environment is a blend of the Mach microkernel and BSD UNIX (which is closer to a monolithic kernel.)

29. Describe how Android uses a unique virtual machine for running Java programs.

Ans: The Dalvik virtual machine is designed specifically for Android and has been optimized for mobile devices with limited memory and CPU processing capabilities.

True/False

30. KDE and GNOME desktops are available under open-source licenses.

Ans: True

31. Many operating system merge I/O devices and files into a combined file because of the similarity of system calls for each.

Ans: True

32. An initial bootstrap program is in the form of random-access memory (RAM).

Ans: False

33. System calls can be run in either user mode or kernel mode.

Ans: False

34. Application programmers typically use an API rather than directly invoking system calls.

Ans: True

35. In general, Windows system calls have longer, more descriptive names and UNIX system calls use shorter, less descriptive names.

Ans: True

36. iOS is open source, Android is closed source.

Ans: False

Chapter 3

Multiple Choice

1. The ____ of a process contains temporary data such as function parameters, return addresses, and local variables.

- A) text section
- B) data section
- C) program counter
- D) stack

Ans: D

2. A process control block ____.

- A) includes information on the process's state
- B) stores the address of the next instruction to be processed by a different process
- C) determines which process is to be executed next
- D) is an example of a process queue

Ans: A

3. The list of processes waiting for a particular I/O device is called a(n) ____.

- A) standby queue
- B) device queue

- C) ready queue
- D) interrupt queue

Ans: B

4. The _____ refers to the number of processes in memory.
- A) process count
 - B) long-term scheduler
 - C) degree of multiprogramming
 - D) CPU scheduler

Ans: C

5. When a child process is created, which of the following is a possibility in terms of the execution or address space of the child process?
- A) The child process runs concurrently with the parent.
 - B) The child process has a new program loaded into it.
 - C) The child is a duplicate of the parent.
 - D) All of the above

Ans: D

6. A _____ saves the state of the currently running process and restores the state of the next process to run.
- A) save-and-restore
 - B) state switch
 - C) context switch
 - D) none of the above

Ans: C

7. A process may transition to the Ready state by which of the following actions?
- A) Completion of an I/O event
 - B) Awaiting its turn on the CPU
 - C) Newly-admitted process
 - D) All of the above

Ans: D

8. In a(n) _____ temporary queue, the sender must always block until the recipient receives the message.
- A) zero capacity
 - B) variable capacity
 - C) bounded capacity
 - D) unbounded capacity

Ans: A

9. A blocking `send()` and blocking `receive()` is known as a(n) _____

- A) synchronized message
- B) rendezvous
- C) blocked message
- D) asynchronous message

Ans: B

10. Which of the following is true in a Mach operating system?

- A) All messages have the same priority.
- B) Multiple messages from the same sender are guaranteed an absolute ordering.
- C) The sending thread must return immediately if a mailbox is full.
- D) It is not designed for distributed systems.

Ans: A

11. When communicating with sockets, a client process initiates a request for a connection and is assigned a port by the host computer. Which of the following would be a valid port assignment for the host computer?

- A) 21
- B) 23
- C) 80
- D) 1625

Ans: D

12. A(n) _____ allows several unrelated processes to use the pipe for communication.

- A) named pipe
- B) anonymous pipe
- C) LIFO
- D) ordinary pipe

Ans: B

13. Which of the following statements is true?

- A) Shared memory is typically faster than message passing.
- B) Message passing is typically faster than shared memory.
- C) Message passing is most useful for exchanging large amounts of data.
- D) Shared memory is far more common in operating systems than message passing.

Ans:A

14. Imagine that a host with IP address 150.55.66.77 wishes to download a file from the web

server at IP address 202.28.15.123. Select a valid socket pair for a connection between this pair of hosts.

- A) 150.55.66.77:80 and 202.28.15.123:80
- B) 150.55.66.77:150 and 202.28.15.123:80
- C) 150.55.66.77:2000 and 202.28.15.123:80
- D) 150.55.66.77:80 and 202.28.15.123:3500

Ans:C

15. Child processes inherit UNIX ordinary pipes from their parent process because:

- A) The pipe is part of the code and children inherit code from their parents.
- B) A pipe is treated as a file descriptor and child processes inherit open file descriptors from their parents.
- C) The STARTUPINFO structure establishes this sharing.
- D) All IPC facilities are shared between the parent and child processes.

Ans:B

16. Which of the following statements is true?

- A) Named pipes do not allow bi-directional communication.
- B) Only the parent and child processes can use named pipes for communication.
- C) Reading and writing to ordinary pipes on both UNIX and Windows systems can be performed like ordinary file I/O.
- D) Named pipes can only be used by communicating processes on the same machine.

Ans: C

17. Which of the following is not a process type in the Chrome browser?

- A) Plug-in
- B) Renderer
- C) Sandbox
- D) Browser

Ans: C

18. The _____ application is the application appearing on the display screen of a mobile device.

- A) main
- B) background
- C) display
- D) foreground

Ans: D

19. A process that has terminated, but whose parent has not yet called wait(), is known as a _____ process.

- A) zombie
- B) orphan
- C) terminated
- D) init

Ans: A

20. The _____ process is assigned as the parent to orphan processes.

- A) zombie
- B) init
- C) main
- D) renderer

Ans: B

Essay

21. Name and describe the different states that a process can exist in at any given time.

Ans: The possible states of a process are: new, running, waiting, ready, and terminated. The process is created while in the new state. In the running or waiting state, the process is executing or waiting for an event to occur, respectively. The ready state occurs when the process is ready and waiting to be assigned to a processor and should not be confused with the waiting state mentioned earlier. After the process is finished executing its code, it enters the termination state.

22. Explain the main differences between a short-term and long-term scheduler.

Ans: The primary distinction between the two schedulers lies in the frequency of execution. The short-term scheduler is designed to frequently select a new process for the CPU, at least once every 100 milliseconds. Because of the short time between executions, the short-term scheduler must be fast. The long-term scheduler executes much less frequently; minutes may separate the creation of one new process and the next. The long-term scheduler controls the degree of multiprogramming. Because of the longer interval between executions, the long-term scheduler can afford to take more time to decide which process should be selected for execution.

23. Explain the difference between an I/O-bound process and a CPU-bound process.

Ans: The differences between the two types of processes stem from the number of I/O requests that the process generates. An I/O-bound process spends more of its time seeking I/O operations than doing computational work. The CPU-bound process infrequently requests I/O operations and spends more of its time performing computational work.

24. Explain the concept of a context switch.

Ans: Whenever the CPU starts executing a new process, the old process's state must be preserved. The context of a process is represented by its process control block. Switching the CPU to another process requires performing a state save of the current process and a state restore of a different process. This task is known as a context switch. When a context switch occurs, the kernel saves the context of the old process in its PCB and loads the saved context of the new process scheduled to run.

25. Explain the fundamental differences between the UNIX `fork()` and Windows `CreateProcess()` functions.

Ans: Each function is used to create a child process. However, `fork()` has no parameters; `CreateProcess()` has ten. Furthermore, whereas the child process created with `fork()` inherits a copy of the address space of its parent, the `CreateProcess()` function requires specifying the address space of the child process.

26. Name the three types of sockets used in Java and the classes that implement them.

Ans: Connection-oriented (TCP) sockets are implemented with the `Socket` class. Connectionless (UDP) sockets use the `DatagramSocket` class. Finally, the `MulticastSocket` class is a subclass of the `DatagramSocket` class. A multicast socket allows data to be sent to multiple recipients.

27. What is a loopback and when is it used?

Ans: A loopback is a special IP address: 127.0.0.1. When a computer refers to IP address 127.0.0.1, it is referring to itself. When using sockets for client/server communication, this mechanism allows a client and server on the same host to communicate using the TCP/IP protocol.

28. Explain the purpose of external data representation (XDR).

Ans: Data can be represented differently on different machine architectures (e.g., *little-endian* vs. *big-endian*). XDR represents data independently of machine architecture. XDR is used when transmitting data between different machines using an RPC.

29. Explain the term marshalling.

Ans: Marshalling involves the packaging of parameters into a form that can be transmitted over the network. When the client invokes a remote procedure, the RPC system calls the appropriate stub, passing it the parameters provided to the remote procedure. This stub locates the port on the server and marshals the parameters. If necessary, return values are passed back to the client using the same technique.

30. Explain the terms "at most once" and "exactly once" and indicate how they relate to remote procedure calls.

Ans: Because a remote procedure call can fail in any number of ways, it is important to be able to handle such errors in the messaging system. The term "at most once" refers to ensuring that the server processes a particular message sent by the client only once and not multiple times. This is implemented by merely checking the timestamp of the message. The term "exactly once" refers to making sure that the message is executed on the server once and only once so that there is a guarantee that the server received and processed the message.

31. Describe two approaches to the binding of client and server ports during RPC calls.

Ans: First, the binding information may be predetermined, in the form of fixed port addresses. At compile time, an RPC call has a fixed port number associated with it. Second, binding can be done dynamically by a rendezvous mechanism. Typically, an operating system provides a rendezvous daemon on a fixed RPC port. A client then sends a message containing the name of the RPC to the rendezvous daemon requesting the port address of the RPC it needs to execute. The port number is returned, and the RPC calls can be sent to that port until the process terminates (or the server crashes).

32. Ordinarily the `exec()` system call follows the `fork()`. Explain what would happen if a programmer were to inadvertently place the call to `exec()` before the call to `fork()`.

Ans: Because `exec()` overwrites the process, we would never reach the call to `fork()` and hence, no new processes would be created. Rather, the program specified in the parameter to `exec()` would be run instead.

33. Explain why Google Chrome uses multiple processes.

Ans: Each website opens up in a separate tab and is represented with a separate renderer process. If that webpage were to crash, only the process representing that the tab would be affected, all other sites (represented as separate tabs/processes) would be unaffected.

34. Describe how UNIX and Linux manage orphan processes.

Ans: If a parent terminates without first calling `wait()`, its children are considered orphan processes. Linux and UNIX assign the `init` process as the new parent of orphan processes and `init` periodically calls `wait()` which allows any resources allocated to terminated processes to be reclaimed by the operating system.

True/False

35. All processes in UNIX first translate to a zombie process upon termination.

Ans: True

36. The difference between a program and a process is that a program is an active entity while a process is a passive entity.

Ans: False

37. The `exec()` system call creates a new process.

Ans: False

38. All access to POSIX shared memory requires a system call.

Ans: False

39. Local Procedure Calls in Windows XP are similar to Remote Procedure Calls.

Ans: True

40. For a single-processor system, there will never be more than one process in the Running state.

Ans: True

41. Shared memory is a more appropriate IPC mechanism than message passing for distributed systems.

Ans: False

42. Ordinary pipes in UNIX require a parent-child relationship between the communicating processes.

Ans: True

43. Ordinary pipes in Windows require a parent-child relationship between the communicating processes.

Ans: True

44. Using a section object to pass messages over a connection port avoids data copying.

Ans: True

45. A socket is identified by an IP address concatenated with a port number.

Ans: True

46. Sockets are considered a high-level communications scheme.

Ans: False

47. The Mach operating system treats system calls with message passing.

Ans: True

48. Named pipes continue to exist in the system after the creating process has terminated.

Ans: True

49. A new browser process is create by the Chrome browser for every new website that is visited.

Ans: False

50. The iOS mobile operating system only supports a limited form of multitasking.

Ans: True