JNTU ONLINE EXAMINATIONS [Mid 2 - ACA]

1. Assume we have a computer where the clock per instruction (cpi) is 1.0 when all memory accesses hit in the cache. The only data accesses are loads and stores and these total 50 % of the instructions. If the miss penalty is 25 clock cycle and miss rate is 2 %,how much faster would the computer be if all the instructions were cache hits.

a. 0.25
b. 0.75
c. 1.5
d. 1.75

2. Which among the following is not a cache optimization technique

a. Reducing the miss penalty
b. reducing the miss rate
c. reducing the miss penalty or miss rate via serialism
d. reducing the time to hit in cache

3. Average memory access time= _ _ _ _ _ _ _ _ _

a. hit time+miss rate*miss penalty
b. miss rate + hit time*miss penalty
c. miss penalty + hit rate * miss rate
d. hit time+hit rate*miss penalty

4. Memory stall cycles= _ _ _ _ _ _ _ _ _ _

a. number of hits*miss penalty
b. number. of misses* miss penalty
c. Number of hits * hit time
d. number of misses * hit time.

5. If the memory hierarchies of the two computers are identical, the cpu with _ _ _ _ _ _ _ _clock rate has a _ _ _ _ _ _ _ _ _ _ number of clock cycles per miss.

a. higher, smaller
b. higher, larger
c. lower, smaller
d. lower, larger

6. Cache memory works on the _ _ _ _ _ _ _

a. Principle of locality
b. working set
c. global memory
d. thumb rule.

7. when the cpu finds a requested data item in cache, it is called a _ _ _ _ _ _ _

a. cache miss
b. cache hit
c. memory miss
d. memory hit

8. The _ _ _ _ _ _ _ the CPI execution, the _ _ _ _ _ _ _ the relative impact of the fixed no of cache miss clock cycles.

a. lower,lower
b. higher,lower
c. lower, higher
d. higher, higher

9. CPU time = _ _ _ _ _ _ _ _

a. cpu execution clock cycles *(memory stall clock cycles+clock cycle time)
b. cpu execution clock cycles +memory stall clock cycles+clock cycle time
c. (cpu execution clock cycles +memory stall clock cycles)*clock cycle time
d. cpu execution clock cycles+ memory stall clock cycles*clock cycle time

10. Miss penalty is measured in _ _ _ _ _ _ _ _

a. cpu clock cycles for a hit
b. cpu clock cycles for a miss
c. cpu clock cycles to execute an instruction
d. cpu clock cycles to access cache.

11. What are the local and global miss rates of second level cache if in 1000 memory references there are 40 misses in the first level cache and 20 misses in the second level cache

a. 50 % and 20 %
b. 2 % and 50 %
c. 50 % and 2 %
d. 20 % and 50 %

12. Depending on the program, a _ _ _ _ _ _ _ _ victim cache might remove one- quarter of the misses in a 4KB direct mapped data cache

a. one entry
b. two entry
c. four entry
d. five entry

13. _ _ _ _ _ _ _ _ _ _ is the no of misses in cache divided by the total no of memory accesses to this cache

a. local miss rate
b. global miss rate
c. local miss penalty
d. global miss penalty

14. Which miss penalty technique is based on the observation that the cpu normally needs just one word of the block at a time

a. multilevel caches
b. critical word first and early restart
c. giving priority to read misses over writes
d. merging write buffers

15. _ _ _ _ _ _ _ _ cache contains only blocks that are discarded from a cache because of a miss

a. write back
b. write through
c. optimized
d. victim

16. Which miss penalty reduction technique ignores the cpu, concentrating on the interface between the cache and main memory

a. multilevel caches
b. critical word first and early restart
c. giving priority to read misses over writes
d. merging write buffers

17. What is the miss penalty of L1 cache

a. hit time L2 + miss rate L1 * miss penalty L2
b. hit time L2 + miss rate L2 *miss penalty L2

c. hit time L1 + miss rate L2 *miss penalty L2
d. hit time L1+ miss rate L1 *miss penalty L2

18. Write buffers hold the update value of a location needed on a _ _ _ _ _ _ _ _
a. write miss
b. read miss
c. write hit
d. read hit

19. _ _ _ _ _ _ _ caches rely of write buffers
a. write back
b. write through
c. optimized
d. victim

20. Which miss penalty technique combines writes to sequential words into a single block to create a more efficient transfer of memory
a. victim caches
b. merging write buffer
c. multi level caches
d. critical word first and early restart

21. _ _ _ _ _ _ _ _ _ optimization reduces misses by improved temporal locality
a. loop interchange
b. blocking
c. pseudo associative cache
d. prediction

22. Which of the following statements is not true about blocking
a. reduces misses via improved temporal locality
b. it can also be used to help register allocation
c. blocking algorithms operate on entire rows or columns of an array
d. it exploits a combination of spatial and temporal locality

23. The latency of the 21264 data cache is _ _ _ _ _ _ clock cycles
a. 1
b. 2
c. 3
d. 4

24. _ _ _ _ _ _ _ _ optimization improves cache performance without affecting the number of instructions executed
a. loop interchange
b. blocking
c. pseudo associative cache
d. prediction

25. Which miss rate reduction technique reduces miss rates without any hardware changes
a. compiler optimizations
b. way prediction and pseudo associative caches
c. large caches
d. higher associativity

26. The very first access to a block cannot be in the cache, so the block must the brought into the cache. These are called _ _ _ _ _ misses
a. collision
b. conflict
c. capacity
d. compulsory

27. _ _ _ _ _ _ _ conflict misses -are due to going from eight way associative to four- way associative
a. Eight way
b. four way
c. two way
d. one way

28. _ _ _ _ _ _ _ _ misses decreases as associativity increases
a. collision
b. capacity
c. compulsory
d. first reference

29. Which miss rate reduction technique is popular in off chip caches
a. large block size
b. larger caches
c. higher associativity
d. compiler optimization

30. If the upper level memory is smaller than what is needed for a program and a significant percentage of the time is spend moving data between two levels in the hierarchy, the memory hierarchy is said to _ _ _ _ _ _ _ _
a. crash
b. trash
c. fault
d. miss

31. What is the size of the page table, given a 32 bit virtual address, 4 KB pages and 4 bytes per page table entry
a. 2MB
b. 4MB
c. 8MB
d. 16MB

32. consider a logical address space of eight pages of 1024 words each, mapped on to a physical memory of 32 frames . How many bits are there in the physical address
a. 10
b. 13
c. 15
d. 18

33. Virtual memory system include a _ _ _ _ _ _ _ _ bit, since the cost of an unnecessary access to the next lower level is so high
a. valid
b. invalid
c. dirty
d. modify

34. The size of the page table is _ _ _ _ _ _ _ the page size
a. directly proportional to
b. inversely proportional to
c. does not depend on

d. equal to

35. Transferring _ _ _ _ _ _ _ _ _ pages to or from secondary storage, possibly over a network, is more efficient than transferring _ _ _ _ _ _ _ _ pages

a. larger, larger

b. larger, smaller

c. smaller , larger

d. smaller, smaller

36. _ _ _ _ _ _ _ _ memory, divided physical memory into blocks and allocates them to different processes

a. primary

b. cache

c. magnetic bubble

d. virtual

37. Which mechanism allows the same program to run in any location in physical memory

a. reloading

b. Relocation

c. memory mapping

d. linking

38. With virtual memory, the cpu produces _ _ _ _ _ _ _ _ addresses

a. virtual

b. physical

c. relocatable

d. linkage

39. To reduce the address translation time , computers use a cache dedicated to these address translations called a _ _ _ _ _

a. TLB

b. paging buffer

c. segmented buffer

d. inverted page table

40. _ _ _ _ _ _ _ occurs in paging

a. internal fragmentation

b. external fragmentation

c. variable size partition

d. memory interleaving

41. The civilian programs are the _ _ _ _ _ _ trusted and hence, have the _ _ _ _ _ _ limited range of accesses

a. most , most

b. least , least

c. most , least

d. least , most

42. A descriptor pointing to a _ _ _ _ _ _ _ _ segment is placed in the global descriptor table, while a descriptor for a _ _ _ _ _ _ _ _ _segment is placed in the local descriptor table.

a. shared, private

b. private, shared

c. shared, public

d. public, shared.

43. _ _ _ _ _ _ _ added to the cpu protection structure expand memory access protection from two levels to many more

a. rings

b. bounds

c. uses

d. locks

44. In alpha memory management, which protection field allows the kernel to read the data within the page

a. valid

b. user read enable

c. kernel read enable

d. kernel write enable

45. Alpha 21264 employs _ _ _ _ _ _ _ _ TLBs to reduce address translation time

a. 1

b. 2

c. 3

d. 4

46. In _ _ _ _ _ _ _ , the cpu and the memory is shared among several interactive uses at the same time, giving the illusion that all users have their own computers

a. multiprogramming

b. timesharing

c. multiprocessors

d. real time systems

47. In alpha memory management, _ _ _ _ _ _ _ _ _ is reserved for the operating system kernel, has uniform protection for the whole space, and does not use memory management

a. seg 0

b. kseg

c. seg 1

d. seg 2.

48. The alpha memory management uses a _ _ _ _ _ _ _ page table to map the address space to keep the size reasonable

a. inverted

b. TLB

c. two level hierarchical

d. three level hierarchical

49. Segment registers in the IA - 32 contains and index to a virtual memory data structure called a _ _ _ _ _ _ _ table

a. page

b. segment

c. mapping

d. descriptor

50. _ _ _ _ _ _ _ field in the segment descriptor of IA-32 specifies the valid opera tions and protection levels for operations that use this segment

a. limit

b. attributes

c. access bit

d. base

51. what fraction of the original computation can be sequential to achieve a speed up of 80 with

100 processors?
a. 0.25
b. 0.5
c. 0.75
d. 1

52. What is the remote request cost for an application running on a 32- processor Multiprocessors, which has a 400ns time to handle reference to a remote memory. The processors clock rate is 1GHz
a. 400 cycles
b. 200 cycles
c. 100 cycles
d. 800 cycles

53. Which of the following is a false statement regarding distributed shared memory multiprocessors.
a. they are also called NUMA's
b. communication occurs through a shared memory
c. It is a cost - effective way to scale the memory bandwidth
d. It increases the latency for accesses to the local memory

54. A multi computer consisting of completely separate computers connected on a local area network are called _ _ _
a. Nodes
b. Workstations
c. Server
d. Cluster

55. What is the interconnection network used in cray T3E multiprocessors?
a. multiple buses
b. farshypercube
c. 2-way 3D torus
d. 8 X 8 Cross bar

56. Vector architectures are the largest class of processors of _ _ _ _ _ _ type
a. SISD
b. SIMD
c. MISD
d. MIMD.

57. In which of the multiprocessors communication of data is done by explicitly passing messages among the processors?
a. message passing multiprocessors
b. distributive shared memory multiprocessors
c. Symmetric shared memory multiprocessors
d. asymmetric multiprocessors

58. Which is the probable architecture for on-chip multiprocessors
a. MPP
b. SMP
c. moderate scale
d. shared virtual memory processors

59. What is the maximum number of processes in sun star fire servers ?
a. 2048
b. 512
c. 64
d. 32

60. What is the typical remote memory access time of HPV series?
a. 500
b. 300
c. 1000
d. 400

61. Which of the following statement is false regarding write-back cache
a. It harder to find the recent value of a data item
b. they use the same snooping scheme both for cache misses and for writes
c. they generate lower requirements for memory bandwidth
d. they are not preferable in multiprocessors.

62. An attempt to write a block in the _ _ _ _ _ _ _ state always generates a miss, even if the block is present in the cache , since the block must be made _ _ _ _ _ _ _ _ _ _ _ _ _
a. shared , Invalid
b. shared ,exclusive
c. Invalid , exclusive
d. exclusive, shared

63. Replication _ _ _ _ _ _ _ _ latency of access and _ _ _ _ _ _ _ _ _ contention for a read shared data item
a. reduces ,increases
b. reduces , reduces
c. increases ,reduces
d. increases, increases

64. Write - back caches generate _ _ _ _ _ _ requirements for memory bandwidth , and slightly _ _ _ _ _ _ _ _ _ _ _ _ _ in complexity
a. lower, decrease
b. greater, decrease
c. lower, increase
d. greater, increase

65. If an operation is done with out intervening operation then the operation are _ _ _ _ _ _ _ _ _ _
a. deadlocked
b. starved
c. atomic
d. non atomic

66. The behavior of reads and writes to the same memory location is defined as _ _ _
a. coherence
b. consistency
c. dependency
d. serialization

67. Multiple writes to the same word with no intervening reads require _ _ _ _ _ _ _ write broadcast in an update protocol
a. Only one
b. One or two
c. atmost one

d. multiple

68. The processors with the sole copy of a cache block is normally called the _ _ _ _ of the cache block.

a. owner

b. user

c. dictator

d. master

69. If the CPU uses a multi level cache with the _ _ _ _ _ property , then every entry in the Primary cache is required to be in the secondary cache

a. exclusion

b. inclusion

c. diffusion

d. shared

70. _ _ _ _ _ _ _ actions introduce the possibility that the protocol can deadlock

a. atomic

b. non atomic

c. illegal

d. sharing

71. Directory requests need to _ _ _ _ _ _ _ _ _ the set shares and also _ _ _ _ _ _ _ the set to perform invalidations

a. update , update

b. update,read

c. read , update

d. update,update

72. Optimization in directory protocols often add complexity by _ _ _ _ _ _ _ _ _ possibility of deadlock and by _ _ _ _ _ _ _ _ _ the types of messages that must be handled.

a. increasing ,decreasing

b. increasing ,increasing

c. decreasing ,increasing

d. decreasing,decreasing

73. What message type and contents are to be transmitted from the local cache to the home director when a processor P has a read miss at address A request data and make P a read sharer

a. read miss A

b. Read miss P,A

c. Write miss A

d. Write miss P,A

74. _ _ _ _ _ _ _ are used to send a value from the home node back to the requesting node

a. Invaliddate

b. Fetch

c. Data value reply

d. Data write back

75. Which directory request sets the shares to the identity of the new owner and the state of the block remains exclusive.

a. Read miss

b. Data write back

c. write miss

d. Fetch

76. cache coherence is an accepted requirement in _ _ _ _ _ multiprocessors

a. small-scale

b. medium-scale

c. large-scale

d. very large scale

77. _ _ _ _ _ _ _ _ keeps the state of every block that may be cached

a. Record

b. File

c. Directory

d. Database

78. In which state exactly one processor has a copy of the cache block and it has written the block , so the memory copy is out of date

a. shared

b. uncached

c. Invalid

d. exclusive

79. _ _ _ _ _ _ _ _ _ _ is the node where the memory location and the directory entry of an address reside.

a. local node

b. home node

c. remote node

d. source node

80. In a directory protocol in which state no processor has a copy of the cache block?

a. shared

b. uncached

c. exclusive

d. Invalidate

81. What is the number of bus transactions required for all n processors to acquire a lock on a variable simultaneously, assuming they are all spinning when the lock is released at time 0.

a. n*n

b. 2n+1

c. 2n

d. n2 +2n

82. Suppose there are 10 processors on a bus that each they to execute a barrier Simultaneously. Determine the number of bus transactions required for all 10 Processors to reach the barrier, be released from the barrier and exit the barrier

a. 204

b. 205

c. 120

d. 102

83. How many bus transactions are needed to have 10 processors lock and unlock the variable using a queuing lock that updates the lock on a miss?

a. n2 + 2n

b. (3n2 + 11n) / 2-1

c. 2n+1

d. 3n-1

84. The cost of building basic synchronization primitives will be too _ _ _ _ _ _ _ _ _ and will _ __ _ _ _ _ _ as the processor count increases

a. high,decrease

b. high,increase

c. low,decrease

d. low,increase

85. Which of the following statements is false about queuing locks

a. Can be implemented either in Hardware of Software

b. If the lock is free, it is simply returned to the processor.

c. They are used to reduce the performance of barrier operation

d. Queuing lock implementation in software assumes a bus-based multiprocessor

86. _ _ _ _ _ _ _ synchronization primitives returns the value of a memory location and automatically increments it

a. exchange

b. test-and-set

c. fetch-and-increment

d. read-and-update

87. If multiprocessors are attempting to get the lock, each will generate the _ _ _ _ _ _

a. Read

b. update

c. set lock

d. write

88. _ _ _ _ _ _ _ _ _ _ is a common technique for reducing contention in shared resources
including access to shared networks and buses

a. Simple barrier

b. exponential back off

c. sense-reversing barrier

d. true based barrier

89. Hardware queuing locks implementation assumes a _ _ _ _ _ _ _ _ _ multiprocessor

a. bus-based

b. directory-based

c. message-passing

d. Distribute shared memory

90. Synchronization problems are quite acute in _ _ _ _ _ _ _ _ _ microprocessors

a. small scale

b. very small scale

c. medium scale

d. large- scale

91. SMT exploits _ _ _ _ _ _ _ _ _ _ _ _ parallelism on a multiple-issue superscalar and hence it is
included in _ _ _ _ _ _ _ _ _ processors targeted at server markets

a. thread-level, low end

b. thread level, high-end

c. Instruction - level, low end

d. Instruction - level, high end

92. Data cache performs slightly _ _ _ _ _ with SMT, while the l2 cache performs slightly _ _ _ _

a. worse ,better

b. better,worse

c. better,better

d. worse,worse

93. Which of the following statement is false about fine-grained multithreading

a. Switches between threads on each instruction, causing the execution of multiple threads to be
interleaved

b. the CPU must be able to switch threads on every clock cycle

c. it hides the throughput losses that arise from both short and long stalls

d. it fastens the execution of the individual threads

94. _ _ _ _ _ provides unique register identifiers , instructions from multiple threads can be mixed in the data path without confusing sources and destinations across the threads

a. Register renaming

b. Register exchange

c. Register reordering

d. Register prefetch

95. _ _ _ _ _ _ _ _ _ _ _ multithreading achieves the improvement in throughput at the cost of
some overhead

a. super scalar

b. coarse - grained

c. fine - grained

d. simultaneous

96. _ _ _ _ _ _ multithreading switches between threads on each instruction causing the
execution of multiple threads to be interleaved

a. super scalar

b. coarse - grained

c. fine - grained

d. simultaneous

97. _ _ _ _ _ _ _ multithreading uses the insight that a dynamically scheduled processor already
has many of the hardware mechanisms needed to support the integrated exploitation of TLP
through multithreading .

a. super scalar

b. coarse - grained

c. fine - grained

d. simultaneous

98. In the _ _ _ _ _ _ _ _ case the interleaving of threads eliminate fully empty slots

a. super scalar

b. coarse - grained

c. fine - grained

d. simultaneous

99. Simulation results have shown that sharing everything is key to maximizing _ _ _ _ _ _ _ _

Performance

a. super scalar
b. coarse - grained
c. fine - grained
d. simultaneous MT

100. The processor configuration for the evaluation of an SMT extension starts with an aggressive superscalar that has roughly _ _ _ _ _ _ _ _ the capacity of existing superscalar processors in 2001.

a. same as
b. double
c. thrice
d. ten times

101. Th advertised average seek time is 5 ms, the transfer rate is 40MB/Sec, it rotates at 10,000 RPM and the controller overhead is 0.1ms. Assume the disk is idle so that there is no queuing delay. What is the average time to read or write a 512 byte sector for a disk?

a. 8.01
b. 8.21
c. 8.11
d. 8.0

102. For flash ,assume it takes 65ns to read 1 byte,1.5µs to write 1 byte and 5ms to erase 4KB.What are the times required to read and write a 64KB block to flash memory.

a. 178.3,4.3
b. 178.3,178.3
c. 4.3,178.3
d. 4.3,4.3

103. _ _ _ _ _ _ _ _ _ _ _ _ device is used to handle the complexities of disconnect / connect and read ahead in magnetic disks.

a. Disk controller
b. array controller
c. DMA controller
d. I/O controller

104. Array density is measured as _ _ _ _ _ _ _ _ _ _ _ _ _ _

a. Tracks/inch on a disk surface*(bits/inch) on a track
b. Tracks/bits on a disk surface*(bits/inch) on a track
c. Bits/inch on a disk surface*(bits/inch) on a track
d. Tracks/inch on a disk array*(bits/inch) on a track

105. _ _ _ _ _ _ _ _ _ tapes limit the speed at which the tapes can spin with out breaking or jamming.

a. Longitudinal
b. near line
c. latitudinal
d. helical scan

106. The disk surface is divided into concentric circle, designated _ _ _ _ _ _ _ _ _ _ _

a. Tracks
b. sectors
c. platter
d. cylinder

107. Rewritable DVD drives cost _ _ _ _ _ _ times as much as DVD-ROM drives

a. 5
b. 10
c. 50
d. 100

108. Low cost VCRS and camcorders make us of _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ tapes

a. Longitudinal
b. near line
c. latitudinal
d. helical scan

109. _ _ _ _ _ _ _ _ tapes mean access to terabytes of information in tens of seconds.

a. longitudinal
b. near line
c. latitudinal
d. helical scan

110. Storage device used for embedded applications is _ _ _ _ _ _ _ _ _

a. Magnetic disk
b. magnetic tape
c. flash memory
d. DVD

111. Which of the following statement is true regarding I/O bus.

a. They are short
b. high speed
c. Have a wide range in the data bandwidth of the devices connected to them
d. Do not follow a bus standard.

112. Which of the following is a false statement regarding the I/O processor

a. They facilitate simultaneous execution of several processes
b. They have dedicated tasks
c. parallelism they enable is very high
d. Doesn't normally change information

113. Which of the following statement is false

a. Bus serves as a shared communication link between the sub systems.
b. Cost of the bus is high
c. The maximum speed of the bus is largely limited by physical factors.
d. Bus creates a communication bottleneck.

114. Which option for a bus among the following leads to high performance.

a. A multiplex address and data lines
b. Narrower
c. No arbitration
d. Synchronous

115. Split-transaction bus has _ _ _ _ _ _ _ _ _ _ band width, but it usually has _ _ _ _ _ _ _ _ _ latency than a bus that is held during the complete transaction

a. Higher, lower
b. lower, higher
c. Higher, higher
d. lower, lower

116. Clock rate of PCI is _ _ _ _ _ _ _ _ _ _
a. 10 MHz
b. 20 MHz
c. 33 or 66MHz
d. up to 100 MHz

117. Which of the following is a serial I/O bus often used in embedded computers.
a. SCSI
b. PCI
c. PCI-X
d. SPI

118. In _ _ _ _ _ _ _ _ _ _ _ _ , portions of the machines address space are assigned to I/O devices.
a. Isolated I/O
b. Memory mapped I/O
c. Programmed I/O
d. Interruptdriven

119. _ _ _ _ _ _ _ _ _ _ _ _ relieves the CPU from waiting for every I/O event, but many CPU cycles are still spent in transferring data
a. Interrupt-driver I/O
b. Programmed I/O
c. DMA
d. I/O mapped I/O

120. Au1000 includes about _ _ _ _ _ _ _ _ _ DMA channels and _ _ _ _ _ _ _ _ I/O device controllers on chip respectively.
a. 1,2
b. 10,20
c. 100,200
d. 50,100

121. The storage overhead of RAID 6 is _ _ _ _ _ _ _ _ _ _ _ _ that of RAID 5
a. Same as
b. twice
c. thrice
d. half

122. Achieving higher _ _ _ _ _ _ _ _ _ _ requires improvement in software quality and software fault tolerance.
a. Reliability
b. Availability
c. Dependability
d. Operability

123. _ _ _ _ _ _ _ _ faults exist for a limited time and are not recurring
a. Hard
b. Transient
c. Intermittent
d. Permanent

124. _ _ _ _ _ _ _ _ _ _ faults have declined due to a decreasing number of chips in systems, reduced power, and fewer connectors
a. Design
b. Operation
c. Environment
d. Hardware

125. Which among the following is not a fault classification in VAX systems
a. Hardware
b. Operating system
c. design fault
d. system management

126. _ _ _ _ _ _ _ _ _ _ automatically forces accesses to several disks.
a. seeking
b. latency
c. tracking
d. striping

127. N devices generally have _ _ _ _ _ _ _ _ _ the reliability of a single device
a. N
b. 1/N
c. 2N
d. N2

128. Non redundant disk array is often called _ _ _ _ _ _ _ _ _ _ _ _
a. Raid 0
b. Raid 2
c. Raid 3
d. Raid 5

129. The distributed parity organization is _ _ _ _ _ _ _ _ _ _ _
a. Raid 0
b. Raid 3
c. Raid 4
d. Raid 5

130. A system _ _ _ _ _ _ _ _ _ occurs when the actual behavior deviates from the specified behavior
a. Fault
b. error
c. failure
d. change

131. Which of the following is not a characteristic of TPC bench mark.
a. Price is included with the benchmark results
b. The data size must scale in size as throughput increases
c. The benchmark results are not audited
d. An independent organization maintains the benchmarks

132. What is the performance metric of complex query OLTP bench mark
a. Transactions per second
b. new order transactions per minute
c. Queries per hour

d. web interactions per second

133. The work-load is gradually _ _ _ _ _ _ _ _ _ _ until the server software is saturated with hits and the response time _ _ _ _ _ _ _ _ _ _ significantly

a. Increases, increases

b. increases, decreases

c. Decreases, increases

d. Decreases, decreases

134. A quad processor running IIS/Windows 2000 is _ _ _ _ _ _ _ _ _ _ a dual processor running TUX/LINUX.

a. Faster than

b. slower than

c. same as

d. almost equal to

135. _ _ _ _ _ _ reconstruction speed implies _ _ _ _ _ _ _ _ _ application performance

a. Increased, increased

b. increased, decreased

c. Decreased, increased

d. Decreased, decreased

136. Transaction processing is concerned with _ _ _ _ _ _ _ _

a. Data rate

b. bit rate

c. I/O rate

d. failure rate

137. What is the bench mark used for evaluating the performance of WWW Servers

a. SPECSFS

b. TPC

c. SPECWEB

d. SPEC89

138. For every 100 NFS operations per second, the capacity must increase by _ _ _ _ _ _ _

a. 1GB

b. 10GB

c. 100GB

d. 1000GB

139. Disk media failures on writes fall into which category of faults?

a. Transient

b. Intermittent

c. Permanent

d. Operation

140. Linux reconstructs _ _ _ _ _ _ _ _ _ and Solaris reconstructs _ _ _ _ _ _ _ _ _ _

a. Slowly, slowly

b. slowly, quickly

c. quickly, quickly

d. quickly, slowly

141. Assuming that the components lifetimes are exponentially distributed and that failures are independent, compute the failure rate of a disk subsystem with the following components and MTTF.

. 10 disks, each rated at 1,000,000 - hours MTTF

. 1 SCSI controller, 200,000 - hours MTTF.

. 1 power supply, 200,000 - hours MTTF

a. 10/1,000,000

b. 12/1,000,000

c. 22/1,000,000

d. 20/1,000,000

142. Assume a disk subsystem with the following components and MTTF

. 10 disks, each rated at 1,000,000 - hours MTTF

. 1 SCSI controller, 200,000 - hours MTTF.

. 1 power supply, 200,000 - hours MTTF

. 1 FAN 200,000 - hours MTTF.

. 1 SCSI cable, 1,000,000 - hours MTTF.

Compute the MTTF of the system as a whole assuming that the age of the component is not important in probability of failure and that failures are independent

a. 2,700,000-hours

b. 2,000,000-hours

c. 192,860 hours

d. 43,500 hours

143. What is the response time of an I/O system with a single disk, if it gets on average 64 I/O requests per second and the average disk service time in 7.8 ms.

a. 7.8ms

b. 15.6ms

c. 14ms

d. 3.9ms

144. Which of the following violates rule of thumb

a. No disk should be used more than 80 % of the time

b. No disk arm should be seeking more than 60 % of the time

c. No disk string should be utilized less than 40 %

d. No I/O bus should be utilized more than 75 %

145. What is the utilization of seek time per disk, if the time of average seek in 5s with 100 IOPS

a. 100 %

b. 50 %

c. 25 %

d. 1 %

146. In which design of I/O system shows the folly of 100 % utilization

a. Response times of the naive cost-preformance design & evaluation

b. Availability if the naove cost-preformance desogn & evaluation

c. more realistic cost performance design and evaluation

d. more realistic design for availability and its evaluation

147. If the OS uses 50,000 CPU instructions for a disk I/O. What is the maximum IOPS for CPU with 2500 MIPS?

a. 50,000 IOPS

b. 25,000 IOPS

c. 10,000 IOPS

d. 1,00,000IOPS

148. Mean-time until data loss (MTDL) increase with _ _ _ _ _ _ _ _ _ disk reliability, and _ _ _ _ _ _ _ _ _ _ MTTR.

a. Increased, reduced
b. increased, increased
c. reduced, reduced
d. reduced, increased

149. In Performance-tuned organization the disk utilization is _ _ _ _ _ _ _ _ _ _ _
a. 100 %
b. 80 %
c. 70 %
d. 60 %

150. What is the disk access latency for performance and availability tuned organization
a. 238ms
b. 40ms
c. 41ms
d. 90ms

151. The loss of signal strength as it passes through a medium, called _ _ _ _ _ _ , limits the length of the fiber
a. Attenuation
b. multipart fading
c. signal to noise ratio
d. interference

152. Suppose you have 25 magnetic tapes, each containing 40GB.Assume that you have enough tape readers to keep any network busy. How long will it take to transmit the data over a distance of 1 Km using cat5 twisted pair wires at 100M bits/sec.
a. 22.8 hrs
b. 2.3hrs
c. 0.9hrs
d. 0.25hrs

153. By limiting the length to 100 meters,``Cat5" wiring can be used for _ _ _ _ _ _ _ _ _
a. 1M bits /sec
b. 10M bits/sec
c. 100M bits/sec
d. 1000M bits/sec

154. Which of the following statement is true about single mode fiber
a. Poor transmitter
b. more reliable
c. More expensive
d. Easy to attach connectors to single mode

155. Level 3 unshielded twisted pairs was good enough for _ _ _ _ _ _ _ _ _ Ethernet
a. 1M bits /sec
b. 10M bits/sec
c. 100M bits/sec
d. 1000M bits/sec

156. _ _ _ _ transmits digital data as pulses of light
a. Fiber optics
b. twisted pairs
c. base band coax
d. broad band coax

157. _ _ _ _ _ _ _ _ was deployed by cable television companies to deliver a higher rate over a few kilometers.
a. coaxial cable
b. cat 5 UTP
c. level 3 UTP
d. fiber optic cable

158. Which among the following is a simplex media
a. "cat3" UTP
b. "cat 5" UTP
c. coaxial cable
d. fiber optics

159. The diameter of multimode fiber is _ _ _ _ _ _ _ _ _ _ _ _
a. 8 to 9 microns
b. 62.5 microns
c. 100 microns
d. 625 microns

160. _ _ _ _ _ _ _ _ _ sends different streams simultaneously on the same fiber using different wavelengths of light.
a. WDM
b. TDM
c. FDM
d. CDM

161. In a failure-intolerant LAN if thetotal intervals and intervals of no failures are 8974 and 8605 respectively, then what percentages of hours a user can`t get his work done
a. 41 %
b. 4.1 %
c. 4.5 %
d. 5 %

162. If the number of failures of 58 desktop computers on a traditional LAN are 654 and that these failures are equally distributed among work stations, then what percentage of hours a user on a workstation can't get his work done.
a. 0.13 %
b. 1.3 %
c. 13 %
d. 4.1 %

163. _ _ _ _ _ _ followed open standards and have less stringent electrical requirements.
a. I/O buses
b. Memory buses
c. USB
d. network interface

164. Software failures occur _ _ _ _ _ _ _ _ _ than Hardware failures
a. Less frequently
b. more frequently
c. Equally
d. much more frequently

165. _ _ _ _ _ _ _ _ _ latency schemes sacrifice fault tolerance
a. low

b. high
c. moderate
d. very high

166. Memory buses provide _ _ _ _ _ _ _ bandwidth and _ _ _ _ _ _ _ latency than I/O buses

a. Higher,lower
b. lower,higher
c. lower,lower
d. Higher,higher

167. Which is the best I/O technique to send large messages?

a. Programmed I/O
b. Interrupt driven I/O
c. DMA
d. Memory-mapped I/O

168. Succesful standards include _ _ _ _ _ _ _ _ cost and stability

a. High
b. Low
c. Moderate
d. Very high

169. The communication system must have mechanisms for _ _ _ _ _ _ _ of a message in case of failure

a. Recovery
b. Diversion
c. destroying
d. Retransmission

170. _ _ _ _ _ _ _ _ _ _ _ have the ability to work around failed nodes and switches

a. SANs
b. MANs
c. LANs
d. WANs

171. A gigabit per second LAN can fully occupy a _ _ _ _ _ _ _ _ GHZ CPU when running TCP/IP

a. 1.0-1.1
b. 0.8-1.0
c. 0.8-0.9
d. 0.9 to 1.0

172. Which of the following is not true regarding SAN

a. the server is like a firewall for the SAN
b. graceful behavior under congestion is critical for SANS
c. SANs appreciate dropping packets during congestion is critical for SANS
d. protocol overhead in much lower for a SAN

173. What is the size of the smallest packet on the Ethernet

a. 64 bytes
b. 128 bytes
c. 256 Bytes
d. 512 Bytes

174. A SAN that tries to optimize based on shorter distance is _ _ _ _ _ _ _ _ _

a. Infiniband
b. Ethernet
c. ATM
d. Ultra band

175. _ _ _ _ _ _ _ _ _ _ _ is used for congestion control in ATM

a. Carrier sense
b. credit based
c. back pressure
d. debit based

176. Ethernet is codified as IEEE standard _ _ _ _ _ _ _ _ _ _ _

a. 802.2
b. 802.3
c. 802.4
d. 802.5

177. Bridges operate at _ _ _ _ _ _ _ _ _ _ layer of the OSI model

a. Physical
b. Data link
c. Network
d. Transport

178. Routers are _ _ _ _ _ _ _ _ _ _ _ than Bridges

a. faster
b. slower
c. of same speed
d. very faster

179. Which among the following operate at the network layer of OSI model

a. hubs
b. routers
c. bridges
d. switches

180. Which among the following is connection oriented

a. Ethernet
b. Infiniband
c. ATM
d. Fast Ethernet

181. Small SMP's with _ _ _ _ _ _ _ _ processors has much better cost-performance than clusters

a. 6-8
b. 4-6
c. 2-4
d. 1-2

182. _ _ _ _ _ _ _ _ server has the goal that a node can fail or be upgraded without bringing down the whole machine.

a. SPEC WEB
b. SPECSFS
c. IBM series 300
d. Sun Fire 6800

183. _ _ _ _ _ _ _ _ _ availability and _ _ _ _ _ _ _ _ _ extensibility make clusters attractive to service provides for the WWW.

a. High, incremental
b. low, incremental

c. High, decremental
d. low, decremental

184. TPC-C cluster scale by a factor of _ _ _ _ _ _ _ _ in price or processors while maintaining respectable cost-performance.
a. 2
b. 4
c. 8
d. 16

185. Clusters are used for the _ _ _ _ _ _ _ _ computers, NUMA the _ _ _ _ _ _ _ _ computers
a. smaller, largest
b. smaller, smaller
c. largest, smaller
d. largest, largest

186. A Sequential program in a cluster of N machines has _ _ _ _ _ _ _ _ _ _ _ _ _ the memory available compared to a sequential program in a shared memory multiprocessor.
a. 1/N
b. N times
c. 2N times
d. 1/(2N)

187. Administering a cluster of N machines is close to the cost of administering
a. N independent machines
b. a single, big machine
c. N/2 independent machinces

d. a single, small machine

188. Clusters are usually connected using the _ _ _ _ _ _ _ _ of the computer
a. Memory bus
b. I/O bus
c. Network interface
d. USB

189. Clusters gets _ _ _ _ _ _ _ _ performance by scaling.
a. low
b. high
c. moderate
d. very low

190. Memory bus has _ _ _ _ bandwidth and _ _ _ latency.
a. lower, lower
b. higher, higher
c. higher, lower
d. lower, higher

191. Astandard VME rack is _ _ _ _ _ _ _ _ _ inches wide and about _ _ _ _ _ _ _ _ feet tall, with a typical depth of _ _ _ _ _ _ _ _ _ inches respectively.
a. 6,19,30
b. 30,6,19
c. 19,30,6
d. 19,6,30

192. Which design cluster example includes the cost of software, the cost of space, some maintenance costs and operator costs.
a. Cost of cluster hardware alternatives with local disk
b. Cost of cluster hardware alternatives with disk over SAN
c. Cost of cluster options that in more realistic
d. Cost of performance of a cluster for transaction processing

193. The uniprocessor clusters costs _ _ _ _ _ _ times the two-way SMP option, and the 8 way SMP cluster costs _ _ _ _ _ _ _ _ _ _ times the two way SMP.
a. 1.6, 1.1
b. 1.1, 1.6
c. 1,6
d. 1,1

194. IBM RAID controller requires _ _ _ _ _ _ _ _ _ _ disk.
a. SCSI
b. IDE/ATA
c. RAMAC
d. FC-AL

195. Smaller computers are generally _ _ _ _ and _ _ _ _ _ _ _ for a given function compared to the larger computers.
a. Costlier, Slower
b. Cheaper, Slower
c. Cheaper , faster
d. costlier, faster

196. What is the size of L2 Cache in xseries370?
a. 256kB
b. 512kB
c. 1024kB
d. 2048kB

197. The database cost is primarily a _ _ _ _ _ _ _ _ _ _ function of the number of processors.
a. non-linear
b. linear
c. ramp
d. step

198. FC-AL can be connected in a loop with up to _ _ _ _ _ devices.
a. 126
b. 127
c. 128
d. 129

199. Collaction rates are _ _ _ _ _ _ _ per unit as space requirements increases
a. costier
b. moderate
c. cheaper
d. much cheaper

200. Collaction Sites are designed assuming no more than _ _ _ _ watts per square foot.
a. 1
b. 10
c. 100
d. 1000