

Introductory material

- Metcalfe's law
- classification: Flynn's and beyond (multiprocessors, multicomputers, ...)
 - SISD, SIMD, MIMD
 - multiprocessors vs. multicomputers (networks of computers)
 - symmetric multiprocessing
 - bus-based multiprocessors
 - cache coherency problems
- snoopy cache
- scalability beyond bus-based multiprocessors: switched multiprocessors
 - Crossbars
 - NUMA
- distributed system
- single system image
- transparency goals issues (very high-level understanding)
 - location
 - migration
 - replication
 - concurrency
- service models
 - centralized
 - client-server (and workstation)
 - peer-to-peer
 - processor pool

Networking

- broadband vs. baseband
- packets
- circuit & packet switched networking
- Understand the concept of protocols and layering
- I will not ask you to enumerate the OSI protocol stack but you should understand the functions of the network, transport, presentation, and application layers.
- ethernet transmission
 - CSMA/CD
 - unreliable, connectionless communication
- clients and servers and services
- transport endpoints (addresses) vs. machine endpoints
- connection-oriented vs. connectionless protocols
- IP (Internet Protocol)
 - Interconnection of networks and routing
 - IP address vs. ethernet address
 - class-based addresses (why did we have classes A, B, and C?)
 - I will not ask about "special" IP addresses (e.g. broadcast)
 - port number: what is its purpose?
 - TCP vs. UDP
- Protocol encapsulation concept
- ethernet device driver: encapsulate IP within an ethernet packet
- what is meant by routing?

- sockets
 - concept
 - operations: create, bind, connect, listen, shutdown, communication (know what they do, not what the parameters are)
- ARP

Quality of Service

- Service metrics
 - Bandwidth
 - Delay
 - Jitter
 - Errors
- Admission control vs. traffic control
- ATM - key points
 - control over quality of service: fixed-size packets, path setup
 - CBR, VBR, ABR
- IP quality of service
 - Nagle's algorithm
 - Differentiated service vs. guaranteed service
 - Flow detection
 - Traffic shaping vs. traffic policing
 - DiffServ
 - RSVP
 - I will not ask about MPLS or 802.1p
 - RTCP and RTP: understand the goals/concept

Naming and binding

- terminology (don't bother with directory service vs. name service)
- understand naming, binding, what a name is, what an address is
- compound name
- I will not ask you to define naming convention or context
- DNS
 - what does DNS do?
 - How do multiple DNS servers together provide domain name service?
 - root name servers
- iterative vs. recursive name resolution

Clock synchronization

- Clock drift and skew, linear compensating function (understand the terms)
- Cristian, Berkeley algorithms
 - Understand the formula for Cristian's and Berkeley
 - Error bounds with Cristian's algorithm. Remember that errors are cumulative.
- NTP
 - NTP Synchronization subnet
 - NTP strata
 - You don't need to memorize the NTP(SNTP) formula but be sure to understand how it gives you the same result as Cristian's

- (don't hesitate to do the algebraic expansion on paper!)
- You don't need to know the symmetric mode NTP algorithms
- You don't need to know how dispersion and jitter are calculated
- You don't need to know the NTP message structure or validation tests

Logical clocks

- Lamport's algorithm
 - goals
 - happened-before relationship
 - partial ordering
 - algorithm
 - deficiencies
 - remedy for generating unique timestamps
- Vector clocks
 - goals
 - algorithm
 - know how to identify concurrent events

Group communication

- I won't ask about anycast
- closed vs. open groups, hierarchical groups
- reliability: atomic, reliable, unreliable
- ordering: global-time ordered, total ordered, causal ordered, sync ordered, FIFO ordered, unordered multicasts
- sending versus delivering messages, hold-back queue
- IP multicasting
 - IP vs ethernet multicast addresses
 - how is ethernet multicast implemented?
 - know that some bits of an IP class D address are mapped onto an ethernet multicast address (you don't have to remember that it's 23 of the 28 bits)
 - hash of address
 - exact match on a small list
 - driver still has to check whether to reject the packet
 - IGMP
 - multicast-aware router
 - understand the basic protocol
 - you don't need to memorize the distinctions between v1, v2, and v3 of IGMP
 - just know that v1 didn't require a client to leave a group

Remote procedure calls

- concept of a remote procedure call (RPC)
- language vs. OS construct
- stub functions
- marshalling
 - data representation
 - explicit versus implicit typing
- semantics of remote procedure calls
 - failure of RPCs
 - at most once and at least once semantics

- Interface definition language (notation) - purpose
- RPC compiler
- Sun (ONC) RPC
 - service registration (portmapper)
 - service lookup, capabilities
- DCE RPC (improvements over SUN RPC)
 - cell directory
 - UUID
 - multi-canonical marshaling
- Microsoft COM+/DCOM/ORPC
 - improvements in RPC over DCE
 - integration with COM
 - surrogate processes
 - distributed garbage collection: reference counting, ping-pong
- CORBA (general concepts, CORBA services)
- Java RMI (general concepts, service registration and lookup, serializable class, remote class)
 - garbage collection in RMI (dirty, clean)
- XML RPC/SOAP (very general concepts)
 - Purpose of WSDL
- Microsoft .NET remoting (general concepts)
 - channels (binary/TCP, SOAP/HTTP/TCP, binary/named pipes)
 - single call, singleton objects, and client-activated objects
 - Leasing distributed garbage collector (LDGC)
- AJAX/XMLHttpRequest, REST (very general concepts)
- I will not ask about the windows communication framework
- We did not cover the enterprise service bus framework, so that will not be on the exam

Scalability

- farm, geoplex (replicated farm)
- service cloning
- service partitioning
- shared disk vs. shared nothing
- active-active vs. active-passive