

Due to Prof. Walker by 5pm on Wednesday 7 December 2005

typed answers preferred

1. In a distributed file system, what are the tradeoffs between client-initiated and server-initiated cache validation?

In client-initiated validation the client sends messages to the server to check validity, but if those checks are done on every access they mitigate the benefits of caching. However, this method does mean the server doesn't have to maintain client state.

In server-initiated validation the server detects potential conflicts and notifies the clients accordingly, which reduces the load on the client but adds overhead to the server and requires it to maintain state info on the clients.

2. Three sender-initiated load distribution algorithms were studied by eager, Lazowska, and Zohorjan, which differ only in their location policy. Would the use of broadcast or multicast (broadcast to a specific subset of nodes) improve any of these three location policies or would that only decrease performance? Explain.

It wouldn't help with the *random* policy, but it could make the polling in the *threshold* policy more efficient if multiple polls are necessary (though at the cost of requiring more other processors to listen to the poll, which would be particularly detrimental in a heavily loaded system). Similarly, it could get the request for queue length to the processors more quickly in the *shortest* policy.

3. How does supporting Processor Consistency improve the performance of a memory system, and what extra burden does it put on the programmer in return for that improved performance?

It allows some memory accesses to be executed in whatever arbitrary order may improve concurrency and performance, and it allows writes done by each particular processor to be pipelined to improve performance. The burden on the programmer is that he or she must understand what operations in the code are guaranteed to occur in order and which are not, and to enforce ordering when necessary using synchronization primitives.

4. What are the main advantages and disadvantages of public key cryptography over private key cryptography?

Public key cryptography avoids private key cryptography's problem of having to distribute the private key over insecure channels. It also provides authentication of the sender in addition to encryption of the message contents, avoiding the need for a separate authentication server. However, public key encryption is more complicated and thus the encryption/decryption time is slower.

5. How does a hard real-time deadline compare to a soft real-time deadline?

Hard real-time deadlines are needed by life-critical applications such as air traffic control, medical systems, and nuclear power plants, where the deadline must be met and a late result is of little or no value.

Soft real-time deadlines are those which are desirable to meet on time, but where late results may be acceptable (though their value may decline as time progresses).