Comments Fact that majority of students chose not to attend lectures was very evident in answers given in the examination.

Question 1
a)For each application the question was looking for comments on why each of the QoS parameters are important or not important.
Many of the answers failed to give reasons why some of the QoS parameters where not important. Ther was also a tendency to claim that internet radio requires high bandwidth, this is not the case. Finally a number of answers failed to mention that web browsing requires reliability.
b)The question gave three mechanisms for separating data and control and for each wanted an indication of how the separation was achieved, an example of its use and a more generic description of what makes it the appropriate approach.
Many of the questions failed to fully answer the question; some failed to cover all mechanisms, while others failed to give an example or indicate the situations to which the mechanism was best suited. Some of the answers failed to give sufficient detail; for example saying that data as embedded in the control, which is just repeating the question, without describing how this occurs.
c)i) The question specifically asked why caching gives scalability. Where scalability is, for example, number of users supported; it is not response time.
Many of the answers described how caching reduces response time; which it does, but was not what the question asked for. The main fact missed by answers was that caching reduces the load on the server that holds the original of the cached data. Thereby allowing that server to effectively support more clients using that data.
c)ii) The question asks for two approaches to keeping cached data up-to-date. For each of these it was seeking a description of the approach, its advantages/disadvantages and an indication of when it was appropriate.
Some answers failed to describe two approaches; a time-to-live approach was the most popular picked. A significant number of answers suggested hashing the original and cached content and comparing the two hashes; which this could work, however answers gave no indication of how to do the compare (the hashes are on different machines). Another popular approach described was for a server to push it content when altered; answers did not indicate how a web server knows where is content is cached. Finally, a number of answers suggested solutions used in hardware with things like write-through caches; network caches are read-only.

Question 2
a)The question was asking for a simple description of the two approaches and their advantages/disadvantages. Generally, the descriptions were okay. However, the advantages/disadvantages lack correctness. A number of answers failed to indicate that asynchronous encryption is significantly more computationally costly that symmetric encryption. A large number of answers suggested that asynchronous encryption is more secure than synchronous encryption; if this was the case then session keys would not use synchronous encryption. A number of answers suggested that ease of key distribution is different between the two approaches. Distributing a shared key confidential is no easier or harder than sending a public key unaltered. Mechanisms exist to do both.
B)The question was asking for a short description of how the three approaches can be used to achieve authentication. As all three involve proving that the participant know something unique to them or shared between participants, the question is asking how this proof is achieved.
Not all answers attempted to describe how all three approaches achieved this proof. Where a description was given, some lacked sufficient detail to see how the proof was achieved. Common mistakes in answers included not encrypting the nonce so that anyone could return the value or the value plus one. Another common mistake was not having an exchange of two messages between the participants; authentication requires an interaction. Some answers described the use of public/private key pairs in three-way and third party approaches; these do not use public/private key pairs. Other answers described the use of session keys in the public/private key approach; these are not used for authentication in this approach.
C)This question was about how digital certificates work and not the process via which they are created.
Many answers described the process via which certificates are created, but not how they allow the distribution of unaltered public keys. One specific point missed was that a certificate is signed using the CA's private key and that by using the CA's public key it is possible to check that the certificate has not been modified. This leads to the need to get the CA's unaltered public key that leans to chains of trust with an assumption for the root of the chains.
Common mistakes included indicating that the certificate is encrypted, not signed, by the CA and that it was signed using the Cas public key, which would allow anyone to alter it and resign it.
D)This question is about keeping a message confidential, proving who sent it and doing this at minimal computational cost.
Many answers suggested various ways of encrypting the whole message using Alice's and Bob's public/private keys. As encrypting using asymmetric keys is computationally costly, none of these fulfil the minimal computation cost part of the question. By not including the creation of a hash of the message contents, many answers failed to include a way to confirm that the message received was the original message. Applying the wrong keys during decryption will produce a result, it may be junk but there is no way to tell this. A hash of the message signed using Alice's private key both confirms that Alice sent the message and that the original message has been received unaltered. Finally, a number of answers suggested the need for a sequence of message to be exchanged between Alice and Bob, this is not necessary; all of the required things can be achieved using a single message from Alice to Bob.

There seemed to be two groups of students taking this exam. One smallish group who knew the material and gave some good answers and a largish group who knew very little. Possibly having not attended much of the course or allocated enough time to it?
Q3.a.i
This was supposed to be a very simple test of your understanding of voice communication. The numbers were deliberately simplified to make calculation easy. My first surprise was how many students clearly have no idea that

the speed of light varies from medium to medium. The question was deliberately ambiguous about the timing delay constraints for normal conversation though this was stated many times during the lectures! Normal conversation fails if the round trip delay (i.e. the time it takes to get from speaker1 to speaker2 and back again) is longer than 300-400ms (0.3 to 0.4 of a second). Either few of you knew this or you chose silly numbers to make your calculation as simple as possible. I was willing to accept any answers between 24,000 and 48,000km as correct.

Q3.a.ii

This was mislabelled on the exam as (i) but did not cause any problems. This produced a range of answers and I was willing to accept well-argued answers for both longer and shorter distances. There were a few very good answers. Most got the 50:50 choice between longer or shorter to be shorter which was the expected answer.

Q.3.a.iii

As expected there were some good answers to this simple part. However, far too many students clearly have little or no idea how TCP works presumably having not learned the lecture notes or read a textbook or tutorial?

Q3.b.i

A diagram very similar to what was hoped for is in the lecture notes and was explained to those present at the lecture. However, most answers left out and did not mention leaky token buckets, a few did not have any sort of priority based queue.

Q3.b.ii

You either knew this or you didn't, it is in your notes. Clearly many students have no idea how quality of service can be implemented and under what circumstances it can be guaranteed. With a little thought you could make up an answer?

Q3.iii

There were a lot of blank answers for this. A few answers that reproduced the diagram that symbolically shows how the leaky token bucket works – how this equates with pseudo code I'm not sure. It is currently beyond AI code generation tools to generate runnable code from such a diagram to my knowledge! The diagram is actually quite a bit more complex than the code needed to implement a leaky token bucket. There were some very good answers. Not surprisingly students' struggled with the need for two separate threads; one to hand out tokens and one to add tokens. The best answers used a simple integer variable for the number of tokens and did not mention queues and buffers. Too many answers were all about the processing of network packets rather than the tokens used to police them.

Q4.a

Most answers seemed to know what ARP is and what it does. A significant number of students had no idea which given that ARP was the 1st COMP28411 laboratory and was then later covered again in lectures was not expected! Quite a few fairly good answers explained what ARP is and roughly how it works but not how it might be adapted for use with transient mobile devices by adapting the caching time though of course this has some undesirable consequences.

Q4.b

Almost all good answers to this. However I was surprised that a few students still did not know what a MAC address is. There was no right or wrong answer. It was all down to how well you argued your case.

Q4.c

A mixture of answers with some reasonable arguments in favour of both possible answers. Subject to the meaning and what is included within a router this was the wrong choice. Generally those thinking routers were in the majority had fairly shallow reasons for this based (I guess) on their experience with home networks.

Q4.d

This was worth a lot of marks and quite a few answers left this part bare or close to bare thus limiting their final mark to be out of 13 rather than 20. This section was taught towards the end of the course when attendance was very low. The marks reflected this with very little knowledge being shown in far too many cases but also some good answers. This was what is called "book work" so there are no excuses for not having read or studied it at least a little. By the way, Ethernet does not have and never has had an acknowledgement frame!

Q4.c

This was expected to be hard but did produce a few good attempts. Worrying how few students know what the term MAC stands for given its meaning is the same when used to refer to a link layer address or as used in the question. SIP of course is an application layer protocol so not usable in the answer. UDP is transport layer so also not usable.