COMPSCI 314 S1C 06 Assignment 3

Answer Key (posted 5 May 06)

Department of Computer Science
The University of Auckland

Posted 12 April 2006
Due 11:59pm, Wednesday 3 May 2006
in https://adb.ec.auckland.ac.nz/adb/

- This assignment will contribute 40/300 = 13.33% to your coursework mark, and 4% to your overall course mark.
- There are 40 marks available on "regular" questions, and 3 bonus marks are available on very difficult questions.
- Bonus marks will not cause any student's total to exceed 40 marks. However bonus marks will improve the total marks of any student who is not awarded all 40 marks from the regular questions.
- No marks will be awarded if you merely state a correct answer. To obtain full credit, your script must clearly explain *why* your answer is correct.
- Plagiarism will not be tolerated. Your explanations must be in your own words.
- If you require additional information in order to answer a problem, you should briefly explain why this information is necessary and why your assumptions about the "missing values" or "missing facts" are reasonable.
- You may submit your assignment either in PDF format (preferred) or in MS Word.

The first questions refer to Figure 3.7 of your textbook, which is reproduced in lecture slide #20 of set #3. In this figure, "RH" is a repeater hub. You should assume the bridging hub uses the routing procedure described on lecture slides #15, #16 and #17 of set #3.

Q1. Consider what would happen if the Ethernet cables leading to stations #1 and #9 of Figure 3.7 were swapped, so that station #1 is connected to the right-hand repeater hub, and station #2 is connected to the left-hand repeater hub. Note: you may safely disconnect and reconnect Ethernet cables without shutting down a station.

a. Immediately after the network connection of station #1 is changed, station #2 attempts to send a frame to station #1. Will this frame cause the bridging hub to change its forwarding database? Will this frame be received by station #1?

[2 marks]

No, the hub is unable to detect that station #1 has moved, so it will not change its forwarding database (1 mark). The hub will not forward this frame to the other subnet, so this frame will not be delivered (1 mark).

b. Immediately after the network connection of station #1 is changed, station #1 attempts to send a frame to station #2. Will this frame cause the bridging hub to change its forwarding database? Will this frame be received by station #2?

[2 marks]

Yes, this frame arrives on the second port of the bridging hub, so the hub will update its forwarding database (1 mark). The hub will forward this frame to the first subnet, so it is reasonable to assume that it will be received by station #2 (1 mark)– however we can't be absolutely certain that station #2 will receive it.

Q2. If the NIC of station #3 in Figure 3.7 is operating in promiscuous mode (so that it receives and buffers all frames appearing on its port), would station #3 be able to eavesdrop on data sent from station #1 to station #2? **[2 marks]**

Yes, a station in promiscuous mode will receive any frame that appears on its subnet, and station #3 is on the same subnet as station #2, so station #3 will receive any frame sent from station #1 to station #2 (2 marks). Note: if the data is encrypted, then station #3 would have to "crack" the encryption before it is effectively eavesdropping on the data sent from station #1 to station #2.

Q3. If the Bridging Hub in Figure 3.7 is replaced by a VLAN switch, and if *none* of the NICs are IEEE802.1Q compliant, would station #1 be able to send data to station #2 without risk of eavesdropping from station #3? Would station #9 also be able to eavesdrop? **[2 marks]**

The IEEE 802.1Q standard is the only VLAN format described in this class. If we assume the NICs are unable to handle the VLAN format of the VLAN switch, then the additional functionality of the switch is useless. Station #1's communication with station #2 can still be eavesdropped by station #3 (1 mark).

Station #9 will not be able to eavesdrop on frames sent from station #1 to station #2, because these frames will not be forwarded to the second subnet, except during a learning phase (1 mark). Note that station #2 must send a frame at least once every few seconds to avoid having station #1's frames "flood" onto station #9's subnet. Flood frames can be eavesdropped.

Note: if a student assumes that the NICs and the VLAN switch are compliant with some VLAN format other than IEEE 802.1Q, then their answers will be the same as in question Q4 below (2 marks).

Q4. If the Bridging Hub in Figure 3.7 is replaced by a VLAN switch, and if *all* of the NICs are IEEE802.1Q compliant, would station #1 be able to send data to station #2 without risk of eavesdropping from station #3? Would station #9 also be able to eavesdrop? **[2 marks]**

If the VLAN switch is configured so that station #1 and station #2 are in their own virtual LAN, and if the NIC in station #3 does not support a promiscuous mode reception of all VLAN frames, then station #3 would be unable to eavesdrop (1 mark). Some students may instead assume that the IEEE 802.1Q NIC in station #3 can be run in promiscuous mode, and this is not an unreasonable assumption. See e.g. http://docs.hp.com/en/T1453-90001/ch02s07.html. If station #3 can receive 802.1Q frames promiscuously, then it can still eavesdrop.

As noted in the sample answer to the previous question, frames sent from station #1 to station #2 will never be repeated onto the subnet containing station #9, so long as station #2's entry in the VLAN switch remains up to date. So station #9 will not be able to eavesdrop (1 mark). Some students may note that additional security against eavesdropping by station #9 would be obtained if the virtual LAN containing stations #1 and #2 is restricted to stations on their subnet, and if the VLAN switch used some method

of building its routing table which doesn't involve flooding frames containing data from station #1, e.g. source routing using discovery frames. See slide #14 of lecture set #3.

Q5. If all the equipment in Figure 3.7 is IPsec-compliant, would station #1 be able to send data to station #2 without risk of eavesdropping from station #3, using AH in transport mode? **[2 marks]**

Data is not encrypted in AH frames, so eavesdropping would still be possible when station #3 is in promiscuous mode (2 marks).

Q6. If all the equipment in Figure 3.7 is IPsec-compliant, would station #1 be able to send data to station #2 without risk of eavesdropping from station #3, using ESP with NULL encryption, in tunnel mode? **[2 marks]**

NULL-mode encryption in ESP leaves the data in plaintext, so station #3 could still eavesdrop (2 marks).

To answer the following questions, you must refer to part 3 of the IEEE 802.3-2002 standard, which is available for download at http://standards.ieee.org/getieee802/download/802.3-2002_part3.pdf. Warning: this is a 3.6 MB document, which will cost you about $0.10 to download on NetAccount during off-peak hours, and about $0.25 during peak hours. So you may wish to save it to a file system, rather than downloading it multiple times.

To receive full marks on these questions, you must support your explanation with one or more direct quotations from IEEE Standard 802.3-2002, giving section and page number(s) for each of your quotations. Your quotations must be accurate, appropriate, and clearly delimited by quotation marks.

Q7. According to your textbook, at page 193, "[t]he choice of 25m [for the limiting length of a drop cable] was rejected by the standards committee as being too small and, after much debate and lobbying, the maximum length of drop cable was set at 200m." Is this an accurate statement about 1000Base-T, as defined in IEEE 802.3-2002? (Hint: see section 40.1 and section 42.) **[3 marks]**

The textbook seems to be inaccurate here, although it may have been an accurate reference to an earlier version of the IEEE 802.3 standard. In Table 42-1, at page 277, we see that the length of a single 1000Base-T Cat-5 segment is limited to 100 metres. From Table 42-2, we learn that a 200m diameter "collision domain" is possible, if we use a repeater and two 100m segments of Cat-5 UTP.

The tables referenced above are for networks designed under "the simpler, but more restrictive rules of Model 1" (page 276). (3 marks)

Note: some students may note that longer distances are allowed under the less restrictive rules of Model 2. For example, we might use DTEs with delays less than 865 bit-times, or we might use cabling with a shorter delay per metre than the maxima listed in Table 42-3. See Note 2 of page 279. So a 200m cable length may in fact be possible for a half-duplex link. However the physical layer of 1000Base-T is intended to support "operation over 100 meters of Category 5 balanced cabling as defined in 40.7" (page 147).

Q8. According to your lecture slide #25 of set #3, in "Gigabit Ethernet ... [the] maximum channel length [is] reduced to 25 metres". Is this an accurate statement about any of the gigabit Ethernet media defined in IEEE 802.3-2002? **[3 marks]**

If we are limited, for some reason, to using only 1000Base-CX shielded jumper cable, then the maximum segment length is indeed only 25m (Table 42-1, page 277) for this

"short haul copper" (page 129). However Table 42-1lists several other cabling options, all of which allow much longer link lengths, so it seems quite inaccurate to say that the Gigabit Ethernet standard limits "channel lengths" to 25m (3 marks).

Q9. Is it possible to use Gigabit Ethernet to communicate over distances longer than 200m? (Hint: look through all of section 42.) **[3 marks]**

Yes, in fact there are several ways to use Gigabit Ethernet over distances longer than 200m. First of all, a single link can be up to 316m long, if optical fibre (1000Base-LX or –SX) is used: see Table 42-1, page 277. Secondly, multiport bridges can be used to join multiple collision domain networks, see page 275. Thirdly, full-duplex links do not require collision detection, so "the maximum link length between DTEs is limited only by the signal transmission characteristics of the specific link" (page 281); for example, a 5000m link length is listed in Table 38.7 of page 108. (3 marks should be awarded to any of these answers, if it is well-supported by reference to the IEEE standard.)

The following questions refer to your textbook Figures 10.2 and 10.3, which are reproduced in lecture slides #14, #15, and #16 of set #4. The Caesar cipher is described nicely in Wikipedia at http://en.wikipedia.org/wiki/Caesar_cipher.

Q10. What key will set the P-box of Figure 10.2.a(i) to a nibble-swap permutation, that is, one in which the 8-bit character 0xXY is mapped to the 8-bit character 0xYX, for any hex digits X and Y? **[2 marks]**

The required key is 56781234, if we assume that the identity permutation has the key 12345678, that is, if we assume that the most-significant input bit is at the top of the diagram (2 marks). If a student assumes that the least-significant bits are at the top of the diagram then the required key is 43218765. If the student does not make it clear how they are numbering the bits of an 8-bit character, then one mark should be deducted.

Q11. Does there exist a key which will set the P-box of Figure 10.2.a(i) to perform a "Caesar cipher" substitution, that is, one in which the 8-bit character 0xXY is mapped to the 8-bit character (0xXY + 5) mod 256, for any hex digits X and Y? **[Bonus: 1 mark]**

No, a P-box can't compute a Caesar cipher. The Caesar encryption of 0x00 is 0x05. No matter how we set the key to the P-box, it will produce an incorrect 0x00 output if given the input 0x00.

Q12. Does there exist a key which will set the product cipher of Figure 10.3 to perform a "Caesar cipher" substitution? **[Bonus: 2 marks]**

This is a very difficult question. The instructor conjectures that the answer is "no" but he doesn't see a short proof of this.

One possible way to decide the question would be to write a computer program that will construct a list of all ciphers C_i that can be implemented by each round (P_8; $4S_2$) of this 3-round substitution-permutation network. Then we could do a brute-force search for integers i, j, k such that the product $C_iC_jC_k$ is the Caesar cipher. If none are found then we would know that the Caesar cipher cannot be implemented by this network. However this computation is infeasible. Each round can implement $8!(4!)^4 \approx 2^{34}$ different ciphers, so our proposed brute-force search of all ciphers which can be implemented in three rounds would require the construction and examination of about 2^{102} ciphers. (2 marks should be awarded to answers that make some valid, non-trivial argument about the ciphers computable by the product cipher of Figure 10.3.)

The following question refers to lecture slides #19 and #20 of set #4.

Q13. According to Moore's law, the price-performance of computer hardware is improved by a factor of two, every eighteen months. Approximately how much would it cost today, to build a machine that can crack a DES key in 56 hours? **[2 marks]**

According to the lecture slide, in July 1998 the 56-hour crack cost $250,000. Eighteen months later, in January 2000, the crack would cost half as much: $125,000. In July 2001, it would cost about $63,000; in January 2003, it would cost about $32,000; in July 2004, it would cost about $16,000; and in January 2006, it would cost about $8000. It is now a few months later than January 2006, and each month the price goes down by a factor of $(1/2)^{1/18}$, so the crack would cost approximately $7000 today (2 marks)

Note: students need not make an exact calculation using logs or fractional powers to gain full marks. Indeed, students using logs or fractional powers are at risk of reporting an answer to an inappropriately high level of precision. The answer should be a round number ($6000, $7000, or $8000), because Moore's law is far from exact. Even so, full marks should be given to an impossibly precise, but reasonably accurate, answer such as $7812.50.

Q14. A simple authentication protocol is illustrated in your textbook's Figure 10.10. This is reproduced on lecture slide #25 of set #4.

a. Would Trudy be able to authenticate successfully as "Alice", by using all three steps of this protocol? **[2 marks]**

By convention, Alice is the initiator of a cryptographic protocol, so the intruder Trudy is attempting to impersonate a client named "Alice". (Students who assume that Trudy is impersonating a server named "Alice" may gain full marks – but only if they make it clear what they have assumed.) In step 1, Trudy sends a message with U = "Alice", with C_p = some public key whose corresponding private key is known to Trudy, and t_c a timestamp. Trudy might, for example, set up a hotmail account under the name Alice@hotmail.com, and then obtain a digital certificate (or an entry in a public key repository) which references this hotmail account as an "identification".

Trudy's first message is encrypted under the server's public key. It thus is identical to a message Alice might have sent, immediately after she obtains a new public key.. In step 2, the server responds to Trudy's message, having no way to discover the fraud except perhaps by comparing the C_p sent in step 1 by Trudy to some other public key used by Alice in some previous transaction with the server. However if the server refuses to accept Trudy's C_p as Alice's public key, then Alice may have difficulty establishing a new keypair when the old one expires. We really need to know more about the public key infrastructure in order to answer this question accurately. However, unless the server is cautious about what C_p keys it accepts, Trudy will successfully authenticate as Alice.

Essentially the same protocol was analysed on Slide 6 of set 5 of the lecture slides.

(2 marks for any answer that shows the student is able to read this cryptographic protocol, and to make relevant statements about what each party learns as a result of running the protocol.)

b. Would Trudy be able to gain service from the Server, by replaying a copy of message #3 that she has copied from a previous, successful, authentication by Alice? **[2 marks]**

This is known as a replay attack. An appropriate defense is for the Server to refuse to deliver service in step 3 on a timestamp t_s which it has already seen. If the server uses this defense then Trudy will fail. (2 marks for any answer that shows the student is able to read this cryptographic protocol, and to make relevant statements about what each party learns as a result of running the protocol.)

Q15. If the Bank of America sends you their public key in an email message, and then sends you another message that is authenticated by this public key, should you follow the instructions in this email to visit a website that will allow you to open a free bank account and have a chance of winning a valuable prize? **[2 marks]**

No! This is called phishing. Security-conscious organizations will not use email to direct you to a website where they will collect information from you. Nor will they deliver their public key to you by email. Hackers currently "phish" for information by sending email that would be impossible for most people to distinguish from email that a reputable (but not security-conscious) bank might send you. In this particular case: how could you distinguish an institution with homepage www.bankofamerica.com from one that has www.bankofamerica.org as its homepage? If you are currently viewing either of these homepages, are you absolutely sure it is what the "real" Bank of America currently intends to have on its homepage?

Q16. To answer the following questions, you should read the Wikipedia article on https (http://en.wikipedia.org/wiki/Https), and you must have access to a web browser with SSL (or TLS) support.

a. Open a connection to a website whose URL has the prefix "https". Inspect the website's digital certificate – you may need to read your browser's helpfile to find out how to do this. Now answer the following questions: What browser are you using, what website are you visiting, and what did find out about this website by inspecting its certificate? **[4 marks]**

I'm using a fully-patched IE 6.0 for XP Service Pack 2, version 6.0.2900.2180.xpsp_sp2_gdr.050301-1519. I'm visiting https://sitekey.bankofamerica.com/sas/signonScreen.do?state=CN.

I couldn't find any information in my browser's helpfiles to tell me how I could inspect the server's certificate. However I happen to know the (apparently undocumented) trick: clicking on the yellow padlock on the bottom edge of the browser window will allow me to inspect the certificate! Here is what I found out:

- ***"Issued to: sitekey.bankofamerica.com"***
- ***"Issued by: www.verisign.com/CPS Incorp.by Ref.LIABILITY LTD.(c) 97 Verisign".***
- ***"Valid from 6/09/2005 to 7/09/2006"***
- ***Clicking on "Issuer Statement" led me to a Verisign website which presented me with a long and legalistic "Relying Party Agreement Version 2.4". This agreement starts with a warning: "YOU MUST READ THIS RELYING PARTY AGREEMENT ("AGREEMENT") BEFORE VALIDATING A VERISIGN TRUST NETWORK^{SM} DIGITAL CERTIFICATE… IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT SUBMIT A QUERY AND***

DO NOT DOWNLOAD, ACCESS, OR USE ANY VERISIGN CRL…" My interpretation of this agreement is that it is intended to limit Verisign's liability "FOR FRAUD OR FOR DEATH OR PERSONAL INJURY CAUSED BY ITS NEGLIGENCE…" The liability proposed by Verisign is USD $100,000.00 for class-3 certificates such as the one I'm analysing. I was unable to determine whether or not I have already violated the terms of the proposed agreement, because my browser might be (for all I know) automatically "validating" Verisign certificates, without giving me a chance to accept or reject the contractual terms on offer by Verisign.

- ***In the second pane of the certificate information subwindow, I am able to see some rather detailed technical information about the certificate, including its serial number and its hashing method. The hashing method is sha1RSA. (By the way, this hashing method has a slight security weakness, which is expected to increase due to Moore's law, see http://it.slashdot.org/article.pl?sid=05/02/19/1424201&tid=93&tid=172&tid=218.)***
- ***This sitekey.bankofamerica.com certificate is signed by a www.verisign.com certificate which was already in my browser's keystore. The www.verisign.com certificate is signed by the "Verisign Class 3 Public Primary CA.***
- ***The bottom line, according to IE 6.0, is that "This certificate is OK."***

Marking notes: one point (up to a maximum of four) for each non-trivial observation made by a student about their browser or the certificate they view.

b. Try to verify the authenticity of the digital certificate you have found, using only the information available to you from the following sources: your textbook, your lectures in this class, the website you are visiting, and the helpfiles of your browser. Now answer the following questions. Were you able to gain confidence in the validity of this certificate? Do you think it is reasonable to expect non-expert users to inspect certificates? **[3 marks]**

I have confidence that this key is registered to the same webarea (sitekey.bankofamerica.com) that I'm viewing, so in this rather restricted sense it is a valid certificate.

I doubt very much that Verisign would issue a certificate to some website with a similar name (e.g. https://www.bankofamerica.org) which is operated by some entity other than the "real" Bank of America.

Amusingly, but not surprisingly, there is an invalid certificate on offer at https://www.bankofamerica.org! When I visit this site I get a Security Alert subwindow from IE 6.0, with the following message:

> ***Information you exchange with this site cannot be viewed or changed by others. However, there is a problem with the site's security certificate.***
>
> ***(tick) The security certificate is from a trusted certifying authority.***
>
> ***(warning) The security certificate has expired or is not yet valid.***
>
> ***(warning) The name on the security certificate is invalid or does not match the name of the site. Do you want to proceed?***

Viewing the certificate I can immediately see two problems: it is issued to "secure.nethollywood.net", and it is valid from 26/07/2005 to 26/08/2005 (apparently this is a US-style date).

I am curious to know what would be served at https://www.bankofamerica.org but I'm cautions about infecting my laptop. So I google first, discovering that Nethollywood is a web-hosting company. I then inspect Google's webcache for www.bankofamerica.org. It would seem that this webarea was set up by someone who, for a while, attempted to mount a class-action lawsuit against the Bank of America.

Since I can't think of any other likely URL for the "real" Bank of America, I am now very confident that I am not being spoofed by the secure website sitekey.bankofamerica.com. (2 marks)

I do not think it likely that anyone other than a security expert would be able to interpret a digital certificate accurately. (1 mark)

Marking notes: one point for each non-trivial observation made by a student about the validity of the certificate, or about the difficulty non-experts would have in interpreting certificates.

Further notes: From http://www.microsoft.com/technet/prodtechnol/winxppro/support/tshtcrl.mspx I gather that IE 6.0 should access a Verisign CA to obtain its latest revocation list, when I click on the yellow padlock symbol. However I don't think my browser actually asked Verisign for a revocation list, because my software firewall (ZoneAlarm) indicates that no network traffic is being sent from my computer when I click on the yellow padlock, or when I navigate within the resulting subwindow containing information about this certificate. However, when I change the default security settings of IE 6.0 to "Check server certificate status", after some delay I get an error message saying "Revocation information for the security certificate for this site is not available. Do you want to proceed?" In this error window I am offered a chance to view the certificate, and when I click this box I get the same certificate-information subwindow that I get when I click on the yellow padlock on the bottom edge of my browser window.

I presume that if the Verisign class-3 certificate security is ever known to be compromised, Verisign (in view of its offered agreement to limit its liability to the not inconsiderable sum of USD$100,000) would ask Microsoft to issue an urgent update to my browser's trusted certificate store.

Google's top-ranked hit for "Bank of America" is www.bankofamerica.com. Googling on "sitekey Bank of America" leads me to some authoritative-looking and interesting news articles on Bank of America's web-security initiatives, e.g. http://news.com.com/2061-10789_3-5723556.html.
