COMPSCI 725

THE UNIVERSITY OF AUCKLAND

SECOND SEMESTER, 2007
Campus: City

COMPUTER SCIENCE
Software Security
(Time allowed: TWO hours)

NOTE: Attempt **ALL** questions in the 12-page script book provided, using approximately **25** words to answer each 5-mark question, **50** words to answer each 10-mark question, and approximately **75** words to answer each 15-mark question. Total possible: **100 marks.**

This is an ungraded sample exam, which should take you about 25 minutes to complete.
Sample answers from students are shown in blue. Instructor's comments are shown in green.

A. The following questions refer to the CPRM system shown in Figure 4 of Myles *et al.*, "Content Protection for Games", *IBM Sys. J. 41:1*, pp. 119-143, 2006. This figure is reproduced below.

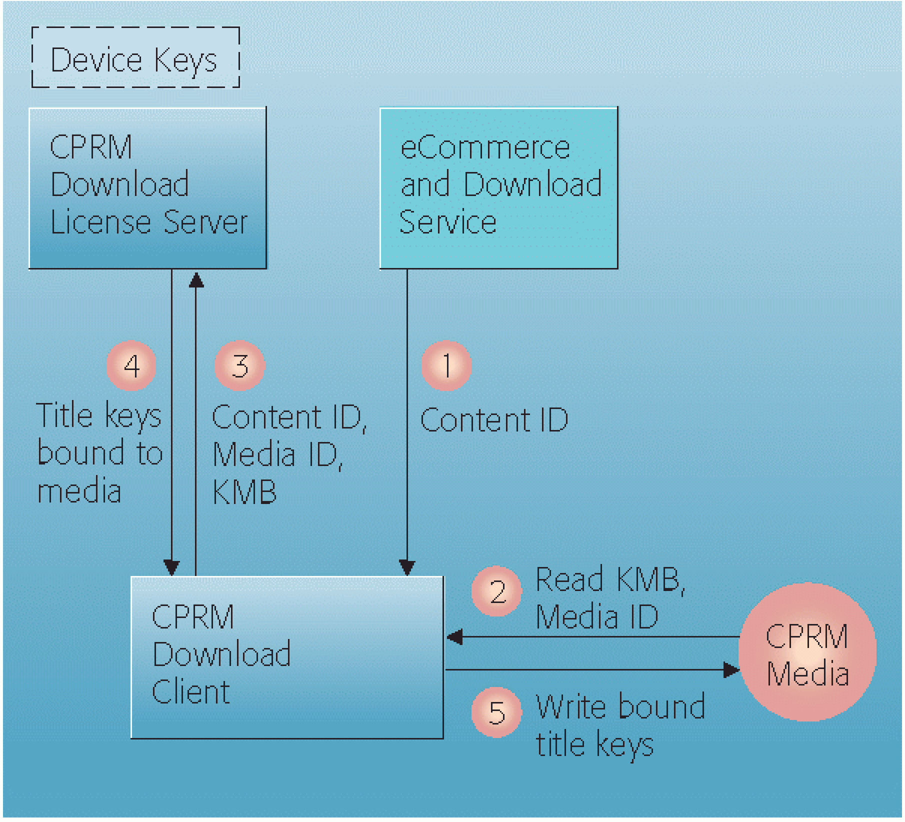
Device Keys
CPRM Download License Server
eCommerce and Download Service
4 Title keys bound to media
3 Content ID, Media ID, KMB
1 Content ID
2 Read KMB, Media ID
CPRM Download Client
CPRM Media
5 Write bound title keys

Figure 4
CPRM for network download

1. Redraw this system diagram, so that it shows an end-user, an attacker, and an author. To obtain full credit, your diagram must clearly show the following

i) the end-user attempting to make an authorised read access (6) to the CPRM-protected media,

ii) the attacker attempting to make an unauthorised read access (7) to this media,

iii) how the author is able to write their intended media content to this protection system. (The arc or arcs for these pre-publication steps should be assigned small numbers, e.g. 0, -1 etc. so that the time-sequence of your diagram is as clear as possible.)

iv) the "trust boundary" in the system, separating the trusted components (and people) from the untrusted components and people, and

v) the valuable item, information, or service which is being protected by this security system. (This should be labelled as "$$".)

Your diagram must be accompanied by a brief explanation of each arc you have added to the original figure 4. **(15 marks)**

Student 1 drew an arc (6) from an End User to eCommerce and Download Service, with the label "Request download with Content (1)". Arc (7) was drawn from Attacker to CPRM Media, with label "Content ID". An unnumbered arc was drawn from the Attacker to the eCommerce and Download Service, with the label "request as user or Author". Arc (8) was drawn from Author to eCommerce and Download Service, with the label "request written writes with Content ID". The Author, Attacker and the End User were all untrusted; all other components were trusted.

The student offered the following explanation.

> **eCommerce & Download Service have the job to Authentication and Authorisation the users. That is, is the user available and what priority he/she have. Normal users request Download with a content ID. The service check his priority for if he can. The answer is yes, process 1-6 go on (6 is the 1st) if not an error message return. An Author does thing same way to get modification permit. An attacker have 2 way to hack the system. 1st is to prtend to a user/author. 2nd is pretend to be an eCommerce and Download service. All parts in original diagram should be trust. But all people shouldn't be trusted.**

This student has not shown much understanding of the CPRM for network download system in their answer. It is possible that an End User could order a download by sending a Content ID, and some payment, to the eCommerce and Download Service. However that step should have been numbered (0) since it would have to precede steps 1 through 5. Furthermore, under this assumption about the e-commerce aspects of this system, the response (1) of the eCommerce and Download Service would have to include, in message (3), a purchase receipt whose integrity can be verified by the CPRM Download License Server.

The student's arc (6) should have been drawn from the CPRM Media to the End User, with a suitable label, such as "Authorised Read Access – the End User's device key, in combination with the Media ID information, is able to decrypt the bound title keys". The student's explanation should have noted that steps 1-5 must have been completed successfully before step 6 can succeeed. After step 5, the user is in possession of a recorded CPRM Media (such as a DVD+R disk), containing some encrypted content and a bound title key.

Part *i:* 1 mark (of 3) to Student 1.

This student's label on their arc (7) doesn't make any sense, for the CPRM Media doesn't use the Content ID. Instead this arc might have been labelled "Unauthorised Read Access – the Attacker's device keys cannot decrypt the bound title keys." The student should also provide an understandable explanation of the attack they have in mind. When I wrote my suggestion for the label on arc (7), I am illustrating a case in which the Attacker has obtained a copy of the CPRM Media created by the End User, and where the Attacker is attempting to play this CPRM Media on some other device. (I might, instead, have assumed that the Attacker is attempting to play the CPRM Media on the End User's own device, in which case the attack would have succeeded – because the CPRM system makes *no* attempt to identify end users. The CPRM Media and the content are identified in this system, but not the End User – it could be a fully anonymous system if the CPRM Download License Server required just a cash payment, but no personal identification, when serving up content in response to message (3).)

The student's other attacking arc, where the Attacker pretends to be an End User or an Author, makes more sense but should have been explained more clearly. For example, if an Attacker obtains the password of an End User who has an account at the eCommerce service, then they *might* be able to convince the service to send them some protected content which is bound to the KMB of the Attacker's media-playing device and to the Media ID of some CPRM-compliant recordable medium which is in the attacker's possession.

Part *ii:* 1 mark (of 3) to Student 1.

The student makes no attempt to describe the steps in the pre-publication process which must occur after an Author identifies and authenticates themselves to the eCommerce system. The student might have shown the Author communicating some content and a suggested retail price (-1) to the CPRM Download License Server, which would assign a new Content ID to this content and then send this Content ID (0) to the eCommerce and Download Service.

Part *iii:* 0 mark (of 3) to Student 1.

The student has drawn a plausible trust boundary, but they have not explained most of their reasoning. Their assertion that "all people shouldn't be trusted" is an over-generalisation. A system is unmaintainable if nobody is trusted enough to make changes to it. A system that lacks a trusted human operator would have to be fully automated, if it is to be functional.

The student has classified the CPRM Download Client as a trusted component, but it is not clear to me why any trust must be placed in this component by the system designers. This component handles encrypted content, and it also handles encrypted key material (such as bound title keys), but it is not given device keys so it cannot decrypt any of this information – and thus it can't violate the confidentiality constraint on the licensed media content. To make a convincing argument in favour of the CPRM Download Client being a trusted component, the student would have to state a security goal that could be violated by an untrustworthy CPRM Download Client.

The student has classified the Author as untrusted. Perhaps the student's intent was to classify nearly everyone (other than the system operators and maintainers who know important secrets such as the device keys) as untrusted "outsiders". If this is the case, and it's not an unreasonable position, then anyone who attempts to gain access as an Author is in the same trust category as a person who attempts to gain access as an End User. However a more careful analysis would indicate that any person who becomes authenticated as an Author will be granted some access rights, in particular they will be able to write content into the system. So the most appropriate

CONTINUED

classification for the Author is, I think, that this actor is trusted. A possible abuse of an Author's write-access would be for them to write content into the system which had been copied (plagiarised) from someone else; such an action could cause harm the true author of this content (who may expect to collect royalties or kudos when End Users obtain licenses to their creations). Here I'm assuming that one of the security goals of the system is to prevent attackers from collecting the royalties (or other benefits) which should be paid to Authors whenever their content (as identified by a Content ID) is purchased by an End User from the CPRM Download License Server. The designers of this system must either trust the Authors not to copy from each other, they will have to include some plagiarism-detection and response mechanism. No such mechanism is shown in this Figure, but since there is no mechanism in this Figure for introducing new content,there is no way to decide whether or not the Authors are trusted. To answer this question we need to whether there are any security goals which are enforced on Authors. If there is any enforcement (prevention; or detection/response) directed at Authors, then they are untrusted – at least with respect to these enforced goals. In the extreme case, if none of the security goals are enforced on Authors (i.e. someone who is authorised to add content into the system), then we would say that the Authors in this system are fully trusted.

The student has correctly classified the trust status of the Attacker. A primary security goal of this system is to restrict read-access to media content. If we trust everyone (who is physically able to access the system) not to make unauthorised read accesses to protected content, then we would not need a technical security mechanism (such as this one) to enforce these restrictions.

The student has correctly classified the trust status of the End User, although their explanation makes me doubt they understand the deeper issues here. In most computer systems, an End User is an authenticated identity. People with End User credentials are generally trusted not to abuse their user privileges, because it would be too expensive and inconvenient (possibly even impossible) to install security mechanisms to prevent all possible abuses of the system by End Users. In the CPRM system shown, there is no authentication of End Users. Anyone who purchases content from the eCommerce and Download Service (step 1) will eventually (in step 5) be provided with content on their CPRM Media device. By this reasoning, the CPRM Media device itself is trusted, but the End Users are not trusted.

The eCommerce and Download Service might accept e-cash payments from unidentified people who want to purchase content. Thus there really need be no authentication of End User identities in the CPRM for network download system.

Student 1 is incorrect to state that the eCommerce and Download Service "have the job to [authenticate] and [authorise] the users". The only entity that is actually being identified and authorised by this system is the End User's CPRM-compliant media storage, recording, and playback device. This apparatus is labelled "CPRM Media" in Figure 4. An identifier (the Media ID) for the CPRM Media and some encrypted key material (in a KMB), are sent in step (3) to the license server, along with the Content ID. In step (4), the licensed media content is issued to the CPRM Download Client. In step (5), the content is written to the CPRM Media.

The CPRM Media is trusted to maintain confidentiality on its device keys. It must also maintain the integrity of its Media ID. Note: the Media ID is a read-only value, essentially a serial number which is unique to each CPRM-compliant recording medium such as a DVD+R disk or the rewriteable storage of an SD card.

If the CPRM Media is ever determined to be untrustworthy (*i.e.* if its device keys are compromised, so that it becomes possible for someone to create a non-compliant CPRM

CONTINUED

playback device with these keys), then no licenses will be issued in the future for the compromised key(s). The existence of such a detection-response regime for CPRM Media implies that this component can become untrusted. So a student who *really* understood the CPRM for network download system would draw the CPRM Media component on the boundary between trusted and untrusted. This hypothetical student would explain that the system generally places trust in the CPRM Media component, but will revoke this trust if all of its keys are compromised. Another possibility for the establishment or withdrawal of the system's trust in the CPRM Media component is explained in the article by Myles *et al.* If the CPRM Media has a connection to a trusted system component, such as the CPRM Download License Server, then the CPRM system can be protected against virtualisation attacks by having the trusted component run a random challenge-response protocol in order to authenticate the CPRM Media. Because this authentication protocol is not shown in Figure 4, it would be most accurate to classify the CPRM Media as a trusted component in that Figure, as has been done by Student 1.

Student 1 has classified the eCommerce and Download Service and the CPRM Download License Server as trusted components. This is an appropriate classification, for these components are trusted to issue access licenses on behalf of Authors (who have previoulsly entrusted their content to this system). Specifically, the CPRM Download License Server is trusted not to breach the confidentiality requirements on the Device Keys and the protected content.

Part *iv:* 2 marks (of 3) to Student 1. This is the fourth part of a very difficult question! The CPRM system is far from easy to understand.

The student has not identified the valuable item for part *v*. To get full credit, they should have put the $$ sign in the CPRM Download and License Server, and they should have explained that this is the portion of the system where Authors (or their publishers) store the content that is being offered for licensed use by End Users. The $$ sign should also be placed in the CPRM media, for this is where a licensed copy of a media title is stored.

Part *v:* 0 marks (of 3) to Student 1.

Student 2 drew an arc (6) from the CPRM Media to the CPRM Download Client. They put a second number (7) on the same arc. This student classified the following as trusted: the Author, the CPRM Download License Server, the eCommerce and Download Service. The End-User, Attacker, CPRM Download Client, and CPRM Media are untrusted. The student placed a valuable item ($$) in the CPRM Download and License Server and in the eCommerce and Download Service. The student drew a circle to show that the End-User and Attacker are near the CPRM Download Client, and their Author was in a circle that was near the CPRM Download License Server. Their explanation reads as follows.

> **(6) An End-user makes an authorised read access. User first do (1) to get content ID by showing his license key. After that he/she can do (2) to get Media ID with KMB. Step (3) send all the information to License server to check both the user and the media are licensed. Then by (4) the user obtain Title keys bound to the media with which the user do (5) to get full access to the CPRM Media.**
>
> **(7) An attacker tries to access the same CPRM Media. If he/she cannot show correct license to the eCommerce Service he cannot get (1). Then the other steps are invalid.**

CONTINUED

Assuming attacker steal a license from someone else and obtain (1) and (20, then step (3) to (5) are all available to the attacker.

This student has demonstrated an excellent level of understanding of the CPRM for network download system. They have made a somewhat surprising assumption (involving a pre-existing "license key") about the previous interaction between the user and the eCommerce and Download service, not shown in this Figure, which resulted in that component sending a Content ID (in step 1) to the user's CPRM Download Client. However this is an examination situation, and I'm impressed that the student has noticed that some message from the user must have provoked the eCommerce and Download Service to issue the Content ID message (1).

Although this student's diagram doesn't clearly show the information flow between the End User and the system, I would give them full marks for parts *i, ii, iv,* and *v.* Because they have neither discussed nor drawn the pre-publication steps I would not give them any credit for part *iii.* Total marks: 12/15 on this very difficult question, this is a very impressive answer! Such high marks on difficult examination questions are (I hope) only achievable by the students who have read the relevant required reading(s) carefully, and who have thought critically and appreciatively about what they have read *before* they sit the examination. The final exam will also contain some easier questions, such as the one below.

2. Butler Lampson, in his article "Computer Security in the Real World", identifies four goals of security: Secrecy, Integrity, Availability, and Accountability. He identifies three basic mechanisms for implementing security: Authentication, Authorisation, and Auditing. His five defensive strategies are isolate, exclude, restrict, recover, and punish. Using Lampson's terminology, identify and briefly describe the most important security goal, mechanism, and defensive strategy of the CPRM system shown in Figure 4. **(10 marks)**

Student 3 responded as follows.

The most important security goal of the CPRM system should be secrecy, because this systems request Media ID and KMB to make sure the CPRM media is the right one and it also get a Content ID from the eCommerce and Download Service to check whether or not the supplied two keys are true. The user would be granted to use the software only if the three keys from different subsystems are valid.

The mechanism of the system is authentication.

The defensive strategy of CPRM should be exclude, because unauthorized user could not identify themselves to the system and could not to access the media.

Student 3 has correctly identified the primary security goal of the system. Their discussion of the keys is reasonably accurate (even though a Media ID is not, strictly speaking, a key). However this discussion is not responsive to the question at hand, which is the definition of the security goal. When defining or discussing a goal, it's inappropriate to discuss how some system might achieve the goal. In Question 1, a system implementation was discussed, and an overall description of this implementation is the focus of the second part (the mechanism) of the current question.

When I find misplaced responses to prior (or future) questions I make some attempt to review my marking of the relevant question. Sometimes a misplaced response will reduce my uncertainty about the student's response to this other question. If I decide that the student

CONTINUED

really didn't understand the issues in that other question, then I'll lower my mark on that question if I had given them the "benefit of the doubt". On the other hand, the misplaced information might convince me that the student *did* really understand how to answer that other question. However… misdirected information in a response suggests that the student does *not* understand the current question, so it will cause me to read the remainder of their response to this question somewhat more critically. In this particular case, I'm satisfied with the explanation of the security goal in the last sentence of the first paragraph, even though it seems a bit odd that the student has assumed that the licensed media content is a software title and not a video or audio title.

Student 3 has asserted that the mechanism is authentication, but they have not explained this mechanism. In their preceding paragraph, the student discussed a system implementation, but they did not clearly distinguish between identifiers (such as a Media ID or a Content ID), authenticators (such a confidential device key), and authorisers (such as a license). Indeed it seems that this student believes the Content ID could be used an authenticator, but it is really only an identification number.

Because this student has confused an identifier with an authenticator, I will not give them high marks on this part of the question.

Student 3 has mounted a reasonable argument in favour of their assertion that the primary defensive strategy is exclusion. I'd give full marks for this part of their answer – although other answers are possible. Perhaps the best possible description of the defensive strategies of this system is that it uses the "restrict" strategy as a backstop to its primary "exclude" strategy. A compromised device key (*i.e.* a failure of the exclusion strategy) could result in previously-issued licensed content becoming available through the use of a "cracked" media player which abuses (relative to the security objectives of the CPRM system) its knowledge of this key. The damage done by such an attack can be restricted, after the crack is detected, by the use of the key-revocation feature of the broadcast encryption algorithm. The system can't prevent attackers continuing to have access to content that is encrypted under the compromised key, but future licensed content can be released under different device keys.

Total marks for Student 3: 7/10. They have shown a good understanding of Figure 4 when describing it, in most cases appropriately, with Lampson's terminology.

Student 4 responded as follows.

> **The keys are the secrecy part which has to be trusted and protected. At the same time they are also the integrity which grants access to a certain media. To ensure accountability and availability all entities have to be present; eCommerce and Download Service and CPRM Download License Server.**
>
> **Furthermore the media is authentication itself by sending its MediaID. Additionally, the KMB has an authenticating function. After every entity (user, media) is authenticated, his authetication has to be send to the CPRM download License Server to get an authorisation (the keys). In cause [case?] of an successful authorisation the accountability of the user to the media is handled with the key. In this system the keys and the media have to be isolated and restricted. The bad guy has to be punished.**

I had great difficulty understanding this student's discussion. From their response to Q1, it is fairly clear that this student is using the vague word "keys" to refer to the "Title keys bound to

CONTINUED

media" which are sent to the CPRM Download Client in step 4. I see no indication that this student understands that there are other keys involved in this protocol, such as the Device Keys shown in Figure 4.

The student has demonstrated some knowledge of the definition of Lampson's terminology, but they have not shown an ability to apply this knowledge in an accurate description of the CPRM for network download system. In particular, keys which are used for decryption are certainly confidential; but the student has described a set of "keys" being sent to the (untrusted) user, saying that these are an "authorisation". There are some keys involved when validating a license using a cryptographic protocol; and a license is an authorisation; so the student has gotten some things right. Total marks: 5/10.

Student 5:

> **Goal: To prevent unauthorised people to have the downloading service and peeking the downloaded source form authorised client.**
>
> **Mechanism: Authentication. Title keys bound to media needs to be written to CPRM Media for client to have the service.**
>
> **Defensive strategies: Restrict. Only client with correct bound title keys can have the downloading service.**

Student 5 has done an adequate job of describing the primary goal, mechanism and defensive strategy of the CPRM for network download system. However they didn't use an appropriate name (secrecy) for the goal, and they have incorrectly used the word "restrict" rather than "exclude" to describe the strategy of preventing an attacker form having access to a system. Perhaps the most worrisome thing about this answer is that it makes me think that the student doesn't understand the system described by Figure 4 – a client doesn't need bound title keys ot access the download service, however if a CPRM Media device doesn't have valid title keys then it will be unable to decrypt and play this title. Total marks: 6/10.

(Other questions…). **[75 marks]**