Remark!: A Secure Protocol for Remote Exams (Transcript of Discussion)

Rosario Giustolisi(✉)

University of Luxembourg, Luxembourg, Luxembourg
rosario.giustolisi@uni.lu

[Omitted explanation of the paper.]

We categorize the security requirements for e-exam in authentication, anonymity and privacy requirements. For instance, we want that only the test answers submitted by registered candidates to be accepted. Similarly, we want that only registered examiners can evaluate the answers submitted by candidates. As an example of an anonymity property, we define anonymous marking, which means that no one can learn the link between a candidate and the answer she submitted. For instance, it is interesting to find out how to guarantee anonymity and authentication properties.

Frank Stajano: It seems to me this set of constraints will severely limit the type of exams you can deliver, because if you have some kind of free-form essay it would be easy to agree on some steganographic hint that, you know, if I am colluding with the examiner then I would say, it's me, if I say the following things.

Reply: Yes, that's true. I think that in remote exams you can communicate to an examiner in the way you said.[1]

Frank Stajano: Right, if I'm going to give you £300,000 if you give me this mark, and then I can always agree, and I send the signal that it was me, if I had a free form. So where does this apply, is this just multiple choice exams?

Reply: The exam setting aims to be as general as possible. You can do a steganographic attack, but it will work if the examiner who marks your test is the one colluding with you. In our protocol, the candidate doesn't know who is going to mark her test. So, the attack can work although the candidate is not 100 % surely it will work.

Alastair Beresford: Presumably if you want the same feature as an old traditional exam, and that you could have collusion, but you want to offer the student the opportunity that they can be anonymous if they want to be. So I could write my own exam, which is made in handwriting, my name is Alistair, please give me 100 % if I wanted to, if I had colluded with Frank for a good mark.

[1] In the post-proceeding version of the paper we clarified that some steganalysis techiniques may help here although threats via subliminal channels are hard to rule out.

B. Christianson et al. (Eds.): Security Protocols 2014, LNCS 8809, pp. 49–54, 2014.
DOI: 10.1007/978-3-319-12400-1_6

Frank Stajano: Yes, you can do it more subtly, so that an auditor won't see that you did.

Alastair Beresford: Sure, and you can do that on physical paper at the moment. So one of the properties I presume you're trying to preserve is the add of a candidate number where the student if they want to be an anonymous candidate.

Frank Stajano: Yes, but what I'm saying is some of these properties are stronger than what we have in a normal exam. So, no one learns the author of a test answer before the mark has been committed. In the current exams thing that we run every year, this can be easily violated. Now, if you are serious about maintaining this then you are really constraining very narrowly the type of exams you can deliver with this.

Joseph Bonneau: It also seems like if it's that constraining then you don't need examiners at all, right, I think there are algorithms could work.

Reply: Yes, if there's an automatic algorithm. Actually, even if it is with open-ended questions I know there is also some algorithm[2], but it depends on the kind of the questions.

Dieter Gollmann: Anonymous examiners, in our examination systems you have the right to come back to me and ask me, why did you mark this question like that? So they have to know who the examiner is. And, if it was somebody else in the university, they would say: it wasn't me so I don't know who marked this.

Peter Ryan: Well, you can then revoke their anonymity.

Reply: Yes, that is also for an anti-revenge aspect.

Dieter Gollmann: In a typical examination situation, students would take a course in software security, and would then know who taught that course in software security, and would know who will have set the examinations on software security, and would know who examined the course for software security. So why try to hide something that is basically obvious?

Reply: It depends. If you think, for instance, a conference, you have anonymous reviewers.

Saar Drimer: But the scenario that you described at the beginning, right, there were brand names of remote exams online based, it doesn't apply to exams in Cambridge, I think, that's my understanding.

James Malcolm: The evaluation can be delegated to some PhD student.

Dieter Gollmann: Well that is different. What in fact is the security requirement?

[2] Two approaches for automatically grading open-ended questions have been proposed in http://aow2012.yolasite.com/resources/aow20120_submission_10.pdf and http://linc.mit.edu/linc2013/proceedings/Session3/Session3Mit-Par.pdf.

Peter Ryan: It may be one way of thinking that is a defence, who's worried about attacking whom, and what things a candidate may be worried about for example, the examiner affecting the results, and this mechanism does at least help that.

Reply: Considering anonymous marking, we want a way to revoke the candidate anonymity because we want to register the mark to someone. We don't want to revoke the anonymity of the examiner to avoid some future possible revenge. If you think about academic conferences you don't know who reviewed your paper.

So, this different anonymity requirements clashes with authentication requirements. There are also secrecy requirements, like question secrecy, to guarantee fairness among the candidates, such as the questions are not revealed before the testing phase. Verifiability properties are also interesting, because the candidate may not need to know who actually marked her test, but know if she got the correct mark for her answer. So, the candidate wants to verify that the examiner will be marking her submitted test, and she also wants to check if she got the mark she deserved.

[Omitted explanation of protocol description.]

Frank Stajano: Is there a different question for every candidate?

Reply: It depends on the exam. If all candidates receive the same question, we lose question privacy, but if they have different questions, or different order of questions, the manager doesn't know to whom each question is given because it's encrypted with the candidate pseudonym.

[Omitted explanation of paper conclusions.]

Yvo Desmedt: Have you taken into account that there are already many, many software systems that actually do these online exams, so for example, at UCL they were using Moodle, which is a system that allows you to basically do these online exams. So when assigning project for students in the classroom of computer security too, I actually did sniff around and found out that this whole communication was basically done in the open, so http was used, not https. So, then two students actually showed to the administration that using an Ipod they could actually change their grade, and only then did the University actually decide to use https instead of http. So there's interest to propose new systems, but what about the security, and the deployment of systems that are already in use.

Reply: So, would you like to use Moodle to run a public competition, or give a job to someone because he passes an exam in Moodle? Or for instance at entrance University exam, in which there is high competition, and you want to allow people from different places to compete...

Yvo Desmedt: I understood the problem, the question is, by proposing something new, there is already a lot of things that have been used, and they may not be good. It's just common.

Peter Ryan: It would be interesting to look at it, yes, we weren't aware of it, so thanks for the tip, and we'll take a look at it. If there are any other systems that

people know of we'd be interested to hear about it. I mean this is not something you could easily stumbled on by Google or whatever.

Yvo Desmedt: No, you look at many universities and see what they actually use. So what seems to me, every different university uses some unique systems, but this is not true, because otherwise there would be as many as there are universities, which is false. But it seems there are many around of these systems. And I don't think many people are looking at how good they are.

Simon Foley: Moodle would be worth looking at. It's entirely Open Source, and it uses a plug-in software architecture for extensions so it might be possible to even implement these algorithms.

Yvo Desmedt: But the thing there was, at UCL, the university decided not to switch the https, they just used open, completely open, so that was an option in Moodle.

Simon Foley: Could these also be seen as examples of potential security vulnerabilities within in Moodle itself? Problems like cross site scripting, and a bunch of other implementation vulnerabilities. While the problem that you're looking at may be different, those types of problems are equally relevant when it comes to implementation.

Peter Ryan: And the other point, so I want to just carry on, what we're trying to stress here is a kind of analogy to what people have to tried to do in e-voting systems trying to minimise the trust, which you have to do, and I think most of these systems would have significant trust in what is the managing system, and we're trying to come up with something which tries to minimise that.

Vashek Matyas: The Blackboard company is developing lot of systems for North American colleges, so you just might, I don't really know whether they denote something like this, but I know that they do have a big coverage of systems there all related, and it's a big company for the area.

Saar Drimer: It seems to me that the main issue here is almost the exam administrators that in academic institutions are interested in this threats that's cheating and what you call plagiarism. Two examples I've seen in the news recently in exams such as TOEFL, and the one equivalent here in the UK. People who were taking the exam would come in, door closes, and then answers being read out. Another example, the people who were taking the exams come in, then the people who were going to take the exam, instead of them coming in and sitting next to them, take the exam, and then go away, that's, they're sitting next to each other for about three hours just reading the books. But that, I can see bits of this system dealing with some of the issues, and make it a bit harder, like randomisation of the questions and answers, and so on, but it doesn't seem to me like it's solving the real issue. Now, the other thing is that there are two different scenarios. There's the university exam taking scenario where there is some form of weak authentication when you go away and they recognise your face. In Cambridge I know that that's kind of system is a weak authentication mechanism.

Joseph Bonneau: But some schools, they are the same in US now.

Saar Drimer: Yes, that's a problem, and then there is the remote, pretty much anonymous, exam taking for getting into university, or getting into a language school, and so on, that here is trumped up in the news quite a bit. So I wonder what bits of this work alleviate some of these concerns.

Reply: Plagiarism, I think, is hard to combat cryptographically in the sense of designing a cryptographic protocol that counters plagiarism problem.

Saar Drimer: Yes, I understand, it's a hard problem.

Reply: So there are some tools like ProctorU and Remote Proctor Now that are designed for invigilation purposes. The UK scandal happened because the invigilator was colluding with all the candidates there, and he read the correct answers aloud.

Joseph Bonneau: So, to extend on Saar's point a little bit, I guess what are the target applications here where people care enough about privacy and having strong security properties is they would do this complex crypto protocol, but they don't have such security concerns that they require candidates to take tests in a secure facility without knowing quite where, etc. I mean, in the US at least it's pretty police state now for the important exams, for admissions, and for like things like I took an IT exam in the fall, and they did like metal detectors to see if I had any devices on me, and then sat down and hardware control, etc. There is like a professional company in the US that does this for multiple different professional exams and things. So there's kind of like that level for tests that people really have jobs riding on. And then it seems like for the Coursera right now that sort of people are mostly taking the courses just for fun really.

Reply: Yes, I think that MOOC plan is to give certification, and you can expect to get credit from the university. So, I think, if we don't propose secure exams with some secure properties, probably they won't use remote exam system to improve their examination system.

Simon Foley: To follow up on that, it might be an interesting to use ceremonies to help model both the physical infrastructure and the system itself.

Peter Ryan: Yes, absolutely. I mean, at the moment these issues are certainly important, but they're just brushed into the assumptions, but I think we will come back to them, and yes, ceremonies why not, maybe the right way to do that.

Joseph Bonneau: Another question. How would this effect the marking process, because, I mean, for smaller scale exams usually it's pretty iterative, where people sit down in a room and they compare answers, and they're sort of developing the criteria while they're looking at the answers, at least in my experience, grading university finals and things like that. I mean, would that violate the examiner anonymity properties here?

Reply: It depends. This is a general approach so if we consider a deterministic marking algorithm, so that the examiner is just a computer process which actually runs the algorithm, we don't have to bother about the examiner identity,

because it's a marking algorithm. It is different when the examiner is a human, for instance with open-ended questions.

Joseph Bonneau: There's a couple of models, like in Cambridge, the one lecturer has to grade everything, right, with no help, right Frank?

Frank Stajano: Yes, well more or less, there may be exceptions, but that's the general model, yes.

Joseph Bonneau: Whereas like in, at least in the US usually like the professor doesn't do the grading, and they make like four TAs sit down in a room and churn through it all, but usually it's pretty collaborative grading as opposed to like one individual person grading each thing.

Daniel Thomas: Well the same thing happens in A-level and GCSE exams in school. The mark scheme that's produced isn't necessarily the mark scheme that's actually used. The examiners read some scripts, look at some answers, then they sit down together and then work out what the actual mark scheme to use is, and then they write that down handwritten on top of the actual mark scheme, and then use that. So again, there the examiners all sit down in a room together and chat before they decide how to mark the questions.

Reply: OK, that's not remote.

Daniel Thomas: That's not remote.

Reply: OK if you focus to this kind of applications, that is, educational. Probably, this is not for public competitions, or job hiring, in which you have not this kind of collaboration. The roles are different, so a candidate cannot participate on mark assignment. Did I understand you correctly?

Daniel Thomas: The examiners sit down together, not the candidates.

Reply: OK. In our protocol we have a pseudonym that belongs to each examiner. So, it would be interesting to see how to face that.