Secure Exams despite Malicious Management

Giampaolo Bella[‡], Rosario Giustolisi[†] and Gabriele Lenzini[†*]
[‡]Dipartimento di Matematica e Informatica
Università di Catania, Italy
[†]Interdisciplinary Centre for Security, Reliability and Trust
University of Luxembourg

***Abstract*—An exam is a practise for assessing the knowledge of a candidate from an examination she takes. Exams are used in various contexts, such as in university tests and public competitions. We begin by identifying various security and privacy requirements that modern exams should meet, especially in the prospect of them being supported by information and communication technologies. These requirements extend well beyond ensuring authenticating the candidate and preventing her from cheating. Cheating is routinely enforced by invigilation by trusted parties, whereas we discuss that an exam should meet its security and privacy requirements against stronger threat models, including malicious exam authorities. Thus exams must be designed with the care normally devoted to security protocols, and in such a mindset we present WATA IV, a new protocol that meets our security and privacy requirements even when an exam manager is malicious.**

I. Introduction

Cambridge Dictionaries Online defines an *exam* as *a test of a student's knowledge or skill in a particular subject that results in a qualification if the student is successful*. The use of exams in meritocratic societies is widespread, with various examples derivable from the education sector (admissions, courseworks and final qualifications) as well as from the work sector (recruitment and progression).

This paper draws its main motivation from the observation that exams raise security and privacy issues that are more challenging than one may think at first. This is due to at least two main reasons. One is that, while threats are traditionally ascribed to candidates and normally mitigated by invigilation, they may as well come from the authorities running the exams. They also may be corrupted to various extents and can, for example, tamper with the tests or alter the marks. Therefore, exams begin to look more balanced in terms of threats or benefits their participants pose or seek. The other reason is that, to increase the security and privacy challenges, it also comes the use of computers. Although the use of information and communication technology simplifies certain tasks occurring during an exam, such as shuffling the questions to put in the tests or recording the marks for the candidates, it generally makes an exam more vulnerable. The traditional strategies to detect and reduce cheating, unless re-designed, may be incapable against digital threats, and the adoption of the new technology must come with new requirements. When overall met, they should ensure at least the same level of security and privacy that paper-and-pencil exams have enjoyed so far. Hence, we shall unfold the argument that an exam must be designed and analysed as carefully as security systems and protocols normally are.

Contribution. It is at least double. It begins with the definition of thirteen security and privacy requirements for exam protocols. The list is foundational by containing the requirements that we deem necessary out of our experience in both running academic examinations and performing security analysis; but, it can be expected that specific exams may demand variations to some requirements or additional ones. In particular, the list includes *anonymous marking*, which safeguards the candidate by stating that her test is anonymous while it is marked. As shall we see, this requirement has been raising interests in the last decade (see §VII).

The second contribution is the computer-assisted exam protocol WATA IV and its informal security analysis. It is the latest in a family of protocols which were incepted back in 2004, with prototypes used at the final exam of the Computer Security module at the University of Catania [1]. Differently its predecessors, WATA IV meets all the thirteen requirements herein identified, and provides a less provincial solution, easily adaptable to different university exam processes. Besides it is robust against a stronger threat model, this including a malicious observer, malicious candidates, an honest-but-curious anonymiser with lightweight participation, and a malicious manager with intensive participation. Such a threat model is realistic in most real situations, even though more malicious situations of theoretical interest can be envisaged.

Outline. After a preliminary and general description of the structure of an exam (§II), this paper identifies the security and privacy requirements (§III), and then gives the threat model and assumptions (§IV) on which WATA IV rests (§V). It continues with the informal analysis of the protocol (§VI), the relevant related work (§VII) and some conclusions (§VIII).

II. Portrait of an exam

Typically, an exam consists of at least four phases: *registration*, which includes preparatory works such as anonymization of the tests; *testing*, which sees candidates take the test; *marking*, when the tests are evaluated; *notification*, when the candidates are notified of their marks.

WATA IV, as explained in §V, is designed to allow the candidate to register for the exam and be notified of her mark at home, while testing takes place traditionally, with the candidate visiting the exam venue.

If a *role* is a set of principals who perform a set of tasks, a typical exam, and WATA IV in particular, insists on the following roles: the *candidate* role, of undertaking the

* G. Lenzini has been supported by FNR-CORE project C11/IS/1183245 "Socio-Technical Analysis of Security and Trust" (STAST).

exam; the *anonymiser* role, of generating the pseudonyms and delivering them at registration; the *manager* role, of handling registration and notification; the *invigilator*, of distributing the tests, supervising the candidates at testing, and gathering the tests; the *examiner* role, of marking the tests. Roles and their participation at the different exam phases are resumed in the following table.

	registration	testing	marking	notification
candidate	×	×		×
manager	×			×
invigilator		×		
anonymiser	×			
examiner			×	

These roles can be taken by various principals, depending on the usage scenarios. At University, the candidate role is played by students, and the examiner and the invigilator role by lecturers; the manager could be any faculty officer, and the anonymiser a third-party service, such as an administrative or secretary office.

III. Security and Privacy Requirements

We identify a number of requirements for exams, which we organized in three lists: authentication, privacy, and other requirements this including secrecy, integrity and verifiability. While we find them highly desirable out of personal experience and discussions with colleagues, the list is not meant to be universal or exhaustive. Certain exams may demand variations, others may demand additional requirements.

Authentication: We identify six authentication requirements pertaining to the candidate, her test, the questions and answers on her test, the mark given to the test, and the notifier.

A1 *Candidate eligibility: if a candidate's identity is entered in the list of candidates registered for an exam, then she passes the official eligibility criteria for the exam.* This requirement insists that only specific candidates can successfully register for that exam.

A2 *Candidate authentication: if a candidate is admitted to testing, then she is is correctly associated with her identity.* This is a rather standard requirement stating that candidates really are who they claim to be.

A3 *Candidate authorisation: if a candidate is admitted to testing, then she is correctly associated with an entry in the list of candidates registered for that exam.*

A4 *Test authentication: if a candidate receives a test from an invigilator, then she is admitted to testing, and **the test** is correctly associated with the candidate identity* This says that a test must be bound to a candidate identity from the moment she receives the test. It also implies that two candidates will be unable to get tested on each other's questions.

We observe the common way to enforce A4 is to assign a test to a candidate only after the candidate inserts in the test the same details that authenticated her. This candidate becomes the test assignee. With exams that are not computer assisted, for example, an authority can check that the candidate writes down the right details on the test sheet, or the authority can write them down personally.

A5 *Answer authentication: if a candidate inserts answers in a test, then she received that test from an invigilator.* This prescribes that a candidate cannot receive a test through cheating and so answer the test of another candidate's, namely a test authenticated by someone else. It prevents candidates to swap their tests when these are already authenticated, ruling out collusion scenarios whereby a knowledgeable candidate agrees to take the test on behalf of a less knowledgeable one, perhaps under compensation.

A6 *Notification request authentication: if the notifier receives a request to attribute a mark to a candidate, then the request is correctly associated with the candidate.* It means that the candidate is the only principal who can decide that she wants to be notified with her mark, and no-one else can do that on her behalf. This is important in scenarios where a candidate is allowed to decide whether to be notified of her mark or not, which interesting pedagogical implications. The candidate might opt not to be notified, for example, having self-assessed her work after testing with an insufficient outcome. Some universities have rules that limits the number of failures during the academic year or force students failing an exam to skip the next exam session. This is also to discourage students who try the exam out without adequate preparation.

It can be observed that the postcondition of certain requirements confirms the precondition of others, producing combined requirements. In particular, the requirement *answer authentication then test authentication then candidate authentication* states that *if a candidate inserts answers in a test, then she is correctly associated with her identity*, and the requirement *answer authentication then test authentication then candidate authorisation then candidate eligibility* states that *if a candidate inserts answers in a test, then she passes the official eligibility criteria for the exam*. Of course, these can be shortened yielding other combined requirements.

Privacy: We identify four privacy requirements pertaining to the test, the mark and the examiner.

P1 *Anonymous marking: The examiner cannot associate a test to its assignee until after he marks (the answer of) the test.* This signifies that the examiner marks a test while ignoring its author: the test is anonymous. It is a clear contribution to the fairness of the marking. As it stands, the requirement insists on anonymity only till the point that the examiner affixes a mark;

P2 *General anonymous marking: No one can associate a test to its assignee until after the examiner marks (the answer of) the test.* This version of anonymous marking generalises the previous one by saying that nobody knows who submitted a test while this is being marked, except the author of the test. An implication is that test anonymity during marking will even resist collusion of the examiner with other authorities.

P3 *Anonymous examiner: The candidate cannot learn the identity of the examiner who marks (the answers of)*

her test. Of course, this requirement can be met only when multiple principals play the examiner role. In particular, it ensures that the test assignee could not try, after the testing but before the examiner in fact marks the test, to coerce or bribe the examiner.

P4 *Mark privacy: if the notifier attributes a mark to a candidate, then this is only revealed to the candidate.* The requirement states that the mark ultimately attributed to a candidate is treated as valuable personal information of the candidate's. It is then up to her to reveal it.

Other requirements: We identify further three requirements of secrecy, integrity and verifiability.

O1 *Question secrecy: the candidate learns the questions on which she will be marked precisely when a test is authenticated, and not before.* This requirement covers both the traditional scenario of questions being kept secret till testing, and that of (typically a high number of) questions being published before testing without indication of which candidate gets which questions at testing. This is a form of temporal secrecy, because it is relaxed when the test begins.

O2 *Test integrity: if a candidate submits a test, then no-one modifies it.* This ensures that each test is passed on to examiners and marked precisely in the form it is submitted, with no modification applicable by anyone to any of its contents, including the questions and the answers.

O3 *Mark verifiability: if the notifier attributes a mark to a candidate, then the candidate can verify that it is the same mark that the examiner gave to the test authenticated by the candidate.* This requirement gives the candidate a sanity check to verify that her mark is correctly attributed. Meeting this in conjunction with anonymous marking would give the candidate a desirable guarantee that the exam has been run fairly.

IV. WATA IV: Threat model and assumptions

The design of WATA IV and its security analysis have been conceived to stand against a specific threat model. This has been set to be as realistic as possible. "Realistic" means that the way in which the attacker threatens the exam should be consistent with situations that one likely expects in a real setting, for example at university exams. A more sophisticated model of attacker with stronger abilities is possible, but considered as future work (see §VI-C).

Threat model

All principals who take the various roles are assumed to be *rational* in the sense that they will not deviate from the protocol unless there is a clear advantage for themselves, in case of collusion, for them and their accomplices. That said, any role can try to achieve malicious goals, and precisely:

Candidates may try to be overmarked, but their interest in colluding depends on the exam scenario. If the purpose of the exam is just to assess their skills, such as with a University exam, a (malicious) candidate is reasonably more inclined to collude with other (malicious) candidates because there may be benefit for all of them. If the exam turns in a competition, such as an admission exam, a (malicious) candidate may not want to collude with others.

Anonymiser is honest-but-curious, hence it is simply interested in collecting the data. Note that the anonymiser is only active at registration, and its being curious captures sufficiently what may happen in real scenarios. There is no real interest in a more active and disruptive maliciousness. For instance, in university exams, the role of anonymiser is often left to administrative or faculty offices. Behaving more maliciously than being curious may trigger controls, which is something that such offices do not want as this is going against their own interests. They are not only accountable for the easy execution of the exam, but any problem will increase their own workload for instance as it happens when the exam has to be repeated. We can see some level of accountability also in public tenders, where third parties can be entitled to be anonymisers. Again, they can be tempted to harvest data for future benefit, but not to play actively against the rules with the risk of compromising their own reputation. Considering more aggressive anonymisers is interesting but out of scope in this paper, and such a study is left for future work (see §VI-C).

Manager and invigilator are malicious in the traditional sense, as they can also perform active attacks. We assume that candidates are invigilated so we exclude forms of collusion between invigilator and candidate roles. However, we observe that such collusion would require a diffuse level of corruption, usually mitigated by admitting many independent invigilators.

In addition to the roles seen above, it is useful to define the *observer* role. This can be seen as a general attacker who does not fit into a canonical exam role. For example, he can watch (parts of) the exam and try to intercept pieces of information. He can harvest private information such as the mark that the examiner ultimately assigns to a candidate, or favour certain candidates. An observer can be an outsider, or a student who is sitting at the exam but only to learn by heart the exam process and the questions for future cheating, a real and serious threat known as *brain dump*.

In particular, WATA IV has been blueprinted to resist the following attacks.

- *The candidate getting a higher mark than her test answers deserve.* This is the most common threat that may come by whoever takes an exam.
- *The honest-but-curious anonymiser.* The anonymiser, which can be played by an internal office or an external service provider such as Google or Microsoft, is curious as it tries to harvest more information than it needs to run its business.
- *The examiner assigning an unfair mark to a specific candidate.* Prejudice and discrimination can bias the evaluation of a test. By learning the author of the test, the examiner may assign a lower or higher mark than the test deserves.
- *The manager tampering with the marks.* The manager wants to notify a candidate with a different mark from the one the examiner assigned to the candidate's test.

This is what happened, for instance, in the Atlanta scandal[1], where teachers and principals changed the marks of all students driven by the illegal objective to increase the ranking of the schools and get more public funds.

- *The manager colluding with the examiner*. The manager may become more effective by colluding with the examiner.

This attack is a refinement of the previous attack. In the Atlanta scandal we have just reported, all authorities colluded to realize the change of the marks.

- *The observer getting candidates' private information.* If he manages to get the marks of various candidates for coercion or trade.
- *The observer tampering with test answers and marks.* For example, an observer wants to downgrade the mark of a candidate to favour another candidate.

Assumptions

We work with a few assumptions.

- *Each candidate has an email address to which principals can send private messages.* We assume that attackers are out of control of the email infrastructure. This means that attacker does not manage the mail account of the candidate.
- *The questions are transmitted securely to the invigilator prior to testing or a very large number of them are made public*. According the exam rules, each test contains one or more questions generated by a question committee. This means that questions are either unknown to a candidate who receives them at testing, or potentially known among a very large set of questions perhaps even before the candidate registers for the exam.
- *A number of principals take the examiner role.* This number should be sufficiently high to make the probability of guessing which examiner marks a specific test as low as the application scenario requires.
- *A number of principals take the candidate role.* This number should be sufficiently high to make the probability of guessing the candidate upon whom a test is authenticated as low as the application scenario requires.

V. WATA IV: Protocol Description

In this section, we first sketch and explain the crypto primitives and the security features that WATA IV adopts, and then present a description of the protocol.

Visual Cryptography: It is a secret sharing scheme devised by Naor and Shamir [2] that allows the human visual system to decrypt a ciphertext split into a number of transparency sheets. The basic version of the scheme is a 2-out-of-2 secret sharing system, where a secret image is split into two images (shares) $share_A$ and $share_B$ which are then printed on transparency sheets. Overlaying the transparency sheets reveals the secret image. The basic scheme is perfectly secure because each share leaks no information about the secret image. In short, the scheme emulates the XOR operation by mapping each pixel of $share_A$ and $share_B$ into a block of 2×2 sub-pixels. Shares $share_A$ and $share_B$ are thus generated so that the result of their overlap is the secret image, whose contrast is reduced by half.

Fig. 1. The candidate paper sheet

QR signature: It generates a QR code that contains a plaintext and its corresponding cryptographic signature. The QR code is a two-dimensional label that facilitates the recognition of data to optical readers. QR codes can store more data in a less space compared to mono-dimensional barcodes. Moreover, they have a customizable level of error correction capability, and can be read from any perspective. Note that the QR code is not a cryptographic primitive: it is a way to represent data. It is worth to mention that QR codes encoding signatures have been already adopted in voting, and precisely by the Prêt à Voter voting system that uses them to enclose the voter choices without including the plaintext [3]. But a self-contained QR signature is more usable for WATA IV, where the plaintext also includes the visual crypto image share.

Commitment Scheme: It is a solution to bind a value to the committer while hiding the value to others. The scheme has two phases: the commitment phase, in which the value is chosen, hidden, and bound to the committer, and the disclosure phase, in which the value is publicly revealed. The Pedersen commitment scheme [4] ensures unconditional hiding, that is, even an unbounded attacker cannot figure out the hidden value. In a nutshell, the Pedersen scheme works as follows. The commitment phase assumes two given public generators $g, h \in \mathbb{G}_q$. The committer chooses the value v and a random commitment parameter $r \in_{\mathcal{R}} \mathbb{Z}_q^*$. The committer then publishes the commitment $c = g^v h^r$. At the disclosure phase, the committer reveals v and r, thus anyone can verify the commitment. WATA IV adopts the Pedersen commitment scheme

[1] e.g., see http://en.wikipedia.org/wiki/Atlanta_Public_Schools_cheating_scandal

during notification. The manager generates a commitment of the candidate's mark. If a candidate decides to be notified and thus reveals her identity to the manager, she can verify that the notified mark is the mark committed by the manager.

The table outlined below resumes the application of the cryptographic primitives according the four phases of an exam.

	registration	testing	marking	notification
Visual Crypto	×	×		×
QR signature	×	×		
Commitment			×	×

A. Protocol Specification

We describe WATA IV through the four exam phases. Figure 3 gives it in form of a message sequence chart.

Registration: The manager creates a new exam identifier, the unique alphanumeric string ex. A candidate who wants to register for the exam uses a private channel to send her personal details (i.e., name, surname, enrolment number and email address) to the manager. We stress that connections to the manager and anonymiser must be through secure channels. The manager checks whether the candidate is eligible for the exam and, if so, enters the candidate details in a dedicated list. After that, the manager forwards the candidate's details and the exam code to the anonymiser.

The anonymiser generates a $token$ that consists of a random alphanumeric string, and a random visual crypto image, $share_A$. Then, the anonymiser generates the second visual crypto image, $share_B$, such that overlapping $share_A$ and $share_B$ results in the image representing of the token, $token_{AB}$. Let $data_A$ denote the triple formed by the candidate's details, ex and $share_A$. The anonymiser signs $data_A$, generates the corresponding QR signature as follows. First, the plaintext $data_A$ is encoded in *Base64* and signed with the signing key of the anonymiser. Then, the signature is also represented in *Base64*, and included in the final QR code with the corresponding encoded plaintext. The anonymiser includes such information in a digital version of an A4 paper sheet $paper_A$ (Figure 1) to facilitate the printing. The QR signatures printed on the bottom of the sheets self contain the data reported on each sheet and its corresponding signature. The anonymiser sends the signed $data_A$, and $paper_A$ to the manager. Similarly, the anonymiser generates $paper_B$ and the signed $data_B$, which include $share_B$ rather than $share_A$. Then, the anonymiser sends the signed $data_B$ and $paper_B$ to the candidate.

For each candidate, the manager stores her corresponding signed $data_A$ into the database, and prints each $paper_A$ on a transparency sheet. Similarly, each candidate prints her $paper_B$ on a common paper sheet, and the registration concludes.

Testing: The manager hands the list of registered candidates, the transparency sheets, and the tests to the invigilator, who brings them at the exam venue. Each candidate takes a seat, and hands a valid identity document to the invigilator for authentication. The invigilator also checks that each candidate is in the list of those registered for the exam. Then, the invigilator finds the transparency sheet that reports the candidate's details, and hands her the transparency sheet along with a test. If some registered candidates fail to show up, some transparency sheets may be undelivered. In that case, the invigilator puts the corresponding transparency sheets and the excess tests aside.

Once the invigilator delivers transparency sheets to all candidates, each candidate can overlay her paper sheet with the corresponding transparency sheet and read the token. The candidate writes down her token in the test sheet and begin to answer the questions. When the testing time is over, the candidate submits her test, and takes the paper and transparency sheets back with her. The candidate can place her test anywhere in the pile of already submitted tests, and this behaves as an anonymous channel. The invigilator collects the pile of tests when all candidates have submitted their tests.

Marking: The invigilator hands all the tests (even the excess) and the remaining transparency sheets to the manager, who, in turn, forwards the filled tests to the examiners. Each examiner chooses a random pile of tests and evaluates all test answers in it, assigns a mark to each test, hence to the token found on the test. The examiner then signs each triple formed by test, token, and mark, and finally sends all such triples to the manager. The latter verifies the signatures and stores the triples in the database. Then, the manager generates a commitment of the each mark, signs each pair of token and mark commitment, and publishes them on a public append-only bulletin board [5]. This allows each candidate to verify whether her test answers were marked while she still ignores how.

Notification: The manager runs the notification for a fixed time frame. The candidate who wants to know her mark sends the signed $data_B$ to the manager via a secure channel. The manager verifies the signature, and overlaps $share_B$ with each $share_A$ into the database until it finds an intelligible token. Notably, this procedure can be implemented, hence require no human involvement. The manager thus retrieves the mark associated with the token from the database, signs them with the corresponding commitment parameter, and sends them to the candidate.

VI. WATA IV: Security Analysis

In the given threat model (§IV) our protocol meets all security requirements seen in Section III, as discussed below.

Authentication

Candidate eligibility (A1) is met because during registration the manager checks that each candidate who wants to register satisfies the relevant criteria.

Candidate authentication (A2) is met because the invigilator checks the identity document of each candidate before admitting her to testing.

Candidate authorisation (A3) is met because the invigilator checks that the identity of each candidate also is in the list of those registered for the exam before the candidate can be admitted to testing.

Test authentication (A4) is met because the invigilator admits a candidate to testing before giving her the transparency sheet that has her name. If a malicious candidate prints a different visual crypto image on her paper sheet, she could then read no intelligible token by overlapping the paper sheet with the transparency sheet. The same applies if any two

malicious candidates swap their paper sheets before testing. As we shall see later, a dispute resolution procedure guarantees that a malicious candidate cannot even claim that no token appears because the manager misprinted the corresponding transparency sheet. Still, a malicious candidate could write a random token on the test, but then at notification the candidate would be unable to send a valid signature of the token and her details ($data_B$).

Answer authentication (A5) is met because the invigilator supervises all through testing, so malicious candidates are unable to swap their tests. However, we are not protected by invigilators that collude with the students, as in the fraud that has been documented by the BBC.[2] In WATA IV we ruled out this attack by assuming independent invigilators to reduce the likelihood of collusions.

Notification request authentication (A6) is met because only the candidate holds her $share_B$.

Privacy

Anonymous marking (P1) is met because the examiner cannot associate a test with a candidate while he is marking it. General anonymous marking (P2) is met too because the manager cannot do that either. Notably, this is the case even if the examiner, the manager, and the invigilator collude, and even after the marking — should a candidate chooses not to get her grade.

Anonymous examiner (P3) is met because an adequate number of principals play the examiner role, and each randomly picks the tests that he will mark. Moreover, the examiner's identity is not revealed to the candidate even after she is notified with the mark.

Mark privacy (P4) is met because the manager notifies the mark to a candidate only if the latter sends a valid $share_B$, and thus each candidate can get only their corresponding mark. Our protocol also guarantees mark privacy against examiners, manager, and invigilator who collude; in fact, they are oblivious to the correct $share_B$. Also, the anonymiser cannot associate a candidate with a mark, because the manager only publishes the commitment of the mark on the bulletin board.

Other requirements

Question secrecy (O1) is met by the assumptions on the origin of tests (§IV) in which the question committee generates the tests.

Test integrity (O2) is met because the invigilator takes a pile of anonymous tests and passes them on to the manager, who in turn hands them to the examiner. None of these principals can benefit from altering the tests, hence they refrain from doing it because they are rational. As we shall see below, in support of this analysis comes the fact that the protocol also meets anonymous marking, consequently mitigating the examiners' interest in tampering with the tests.

Mark verifiability (O3) is met because the candidates can verify the examiner signature when the manager notifies their marks. The protocol meets this requirement even if the manager and the invigilator collude with the examiners, because the manager publicly commits to the marks before being able to associate them with the candidates' identities.

[2]see "Student visa system fraud exposed in BBC investigation" at http://www.bbc.com/news/uk-26024375.

A. Dispute resolution

A notable feature of WATA IV is the support for dispute resolution during testing. In fact, the combination of QR signatures and visual cryptography guarantees an easy procedure to find the culprit if the candidate or the manager misbehave. Such a dispute originates if no intelligible token can be read when the candidate overlaps the paper sheet with the transparency sheet. Should such a dispute arise, the invigilator could then quickly resolve it as follows. He can easily scan the QR code printed on the candidate's paper sheet and check the correctness of the QR signature. Then, he checks if the candidate details revealed by the QR code match the ones written on the candidate's paper sheet. If so, the invigilator overlaps the visual crypto image unveiled by the QR code with the transparency sheet provided by the manager. If this reveals no intelligible token then the manager misprinted the corresponding transparency sheet, otherwise the candidate misprinted her visual crypto image. The outcome of the dispute can be double checked by repeating the procedure with the QR code printed on the transparency sheet.

B. Comparison with predecessor

WATA IV brings along significant security improvements compared to its predecessor WATA III [1]. In particular it is augmented with Mark privacy, Mark verifiability, and General anonymous marking. Moreover, (a) WATA IV meets the security requirements despite a more realistic threat model that sees the manager being actively malicious, and not only honest-but-curious as in WATA III; (b) the candidate receives her visual crypto image directly rather than indirectly through the manager at the exam venue, as instead she does in WATA III. This change impedes that the manager records the candidate's share; (c) the candidate can record the manager's transparency sheet without raising threats. This is possible because tokens are signed, hence cannot be forged to associate a test answer to a candidate other than its author; (d) anyone, even a malicious manager, can register the mark after the notification thanks to Mark Verifiability; (e) any dispute between the candidate and the manager can be solved with no efforts at testing: the candidate makes sure she received a valid transparency sheet if an intelligible token appears. If no token appears, the QR signatures on the sheets reveal who misbehaved.

C. Extension

WATA IV assumes an honest-but-curious anonymiser, a choice which we motivated in (§IV). However, were one interested to discuss an anonymiser that acts maliciously and not only curiously, WATA IV can be extended to preserve its security properties even under this stronger threat model. We envisage such extension by using a 1-out-of-N oblivious transfer protocol (e.g., see [6]), and perhaps an oblivious printing of the shares as suggested in [7]. Therefore, the candidate and the manager jointly generate their shares at registration: neither the candidate nor the manager know each other shares, but still they have guarantees to hold the correct ones. The details of this extension are left for future work.

VII. Related Work

The topic of secure electronic exam is only recently attracting the attention of the scientific community. There are still a limited numbers of works addressing this topic, some of them indirectly. Foley et al. [8] describes confidentiality requirements for Computer Supported Collaborative Working (CSCW). They advance exams as a case study but not to the point to provide a design of a protocol. Other papers study the security requirements for e-learning systems [9], [10]. They focus on authentication and privacy issues of teachers and students, and sketch some ideas towards security [11]. However, while these works focus on external threats, this paper has shown that exams have to cope with various forms of insider threats, such as cheating candidates or corrupt examiners.

Some popular exam systems are adopted in university tests. In INFOSAFE [12] candidates write their personal details on top of a tamper-evident paper, hide them with a flap which is bent and glued over, flap that is tore off after the marking. NEMO-SCAN [13] uses a patented anonymity paper cover that hides the candidate details, which are written aside, while leaving free a section where examiners can mark. The part with the candidate's details is scanned using a proprietary device, which reveals the candidate and assigns her the mark. Other universities, such as Dublin City University and the University of Sheffield, use their own exam systems [14], [15]. However, all such solution assume a trusted manager, contrarily to WATA IV, to achieve anonymous marking, and it remains unclear whether they could scale up to other security properties.

Huszti et al. [16] propose an Internet based exam without trusted third parties, but with an honest role, the Registry. The protocol aims at a small number of security requirements. Arapinis et al. [17] propose a cloud-based conference management system that guarantees secrecy and privacy properties. Their work addresses threats from a malicious cloud by means of an honest manager. Other works focus on countermeasures to avoid plagiarism during on-line exams [18], [19], [20], and advance methods to ensure the authentication of the candidates [21], [22]: they assume trusted managers and ignore privacy properties such as anonymous marking.

VIII. Conclusions

Exam security has a major role on the widespread acceptance of computer assisted protocols for exam competitions. This paper first has identified thirteen foundational security and privacy requirements for exams. Then, it has described WATA IV, an exam protocol that allows remote registration and notification (e.g., at home), but which requires the exam to be taken *in situ*. WATA IV provides the same functional requirements of traditional exams while guaranteeing a number of security properties in a realistic threat model. Its informal analysis confirms that the protocol meets all thirteen requirements. The next step is to analyse WATA IV formally and to investigate how to counter further threats. For example, how to avoid or detect candidate cheating during testing is an interesting topic for a future research after the recent system fraud scandals. We envisage that solutions such as data mining used to derive patterns [23] can be adopted on top of WATA IV to even detect plagiarism threats.

References

[1] G. Bella, G. Costantino, L. Coles-Kemp, and S. Riccobene, "Remote management of face-to-face written authenticated though anonymous exams." in *CSEDU*. SciTePress, 2011, pp. 431–437.

[2] M. Naor and A. Shamir, "Visual cryptography," in *EUROCRYPT*, 1994.

[3] C. Burton, C. Culnane, J. Heather, T. Peacock, P. Y. A. Ryan, S. Schneider, S. Srinivasan, V. Teague, R. Wen, and Z. Xia, "Using pret a voter in victorian state elections," in *Proc. of the 2012 Conference on Electronic Voting Technology/Workshop on Trustworthy Elections*, ser. EVT/WOTE'12. USENIX Association, 2012.

[4] T. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *CRYPTO*, ser. LNCS, J. Feigenbaum, Ed. Springer, 1992, pp. 129–140.

[5] J. Heather and D. Lundin, "The append-only web bulletin board," in *Formal Aspects in Security and Trust*, ser. LNCS, P. Degano, J. Guttman, and F. Martinelli, Eds. Springer, 2009, vol. 5491, pp. 242–256.

[6] W.-G. Tzeng, "Efficient 1-out-of-n oblivious transfer schemes with universally usable parameters," *IEEE Transactions on Computers*, vol. 53, no. 2, pp. 232–240, 2004.

[7] A. Essex, J. Clark, U. Hengartner, and C. Adams, "How to print a secret," in *Proc. of the 4th USENIX Conference on Hot Topics in Security*, ser. HotSec. USENIX Association, 2009.

[8] S. N. Foley and J. L. Jacob, "Specifying security for computer supported collaborative working," *Journal of Computer Security*, vol. 3, pp. 233–253, 1995.

[9] K. El-Khatib, L. Korba, Y. Xu, and G. Yee, "Privacy and security in e-learning," *International Journal of Distance Education*, vol. 1(4), 2003.

[10] E. R. Weippl, "Security in e-learning," *eLearn*, 2005.

[11] G. Kambourakis, D. Kontoni, A. Rouskas, and S. Gritzalis, "A PKI approach for deploying modern secure distributed e-learning and m-learning environments," *Computers & Education*, 2007.

[12] "INFOSAFE," http://www.anonymousmarking.com/.

[13] "NEMOSCAN," http://www.neoptec.com/en/products/technologies-products.php.

[14] "Exam procedures Dublin City University," http://www4.dcu.ie/iss/am/index.shtml.

[15] "Exam procedures at the University of Sheffield," http://www.shef.ac.uk/lets/design/handbook/28.

[16] A. Huszti and A. Petho, "A secure electronic exam system," *Publicationes Mathematicae Debrecen*, vol. 77, no. 3-4, pp. 299–312, 2010.

[17] M. Arapinis, S. Bursuc, and M. Ryan, "Privacy-supporting cloud computing by in-browser key translation." *Journal of Computer Security*, vol. 21, no. 6, pp. 847–880, 2013.

[18] Y. Sabbah, I. Saroit, and A. Kotb, "An interactive and secure e-examination unit (iseeu)," in *Proc. of the 10th RoEduNet Int. Conference, 2011*, 2011, pp. 1–5.

[19] J. Kasprzak and M. Nixon, "Cheating in cyberspace: Maintaining quality in online education," in *Association for the Advancement of Computing In Education*, 2004.

[20] M. Kikelomo, G. Wills, and D. Argles, "User security issues in summative e-assessment security," in *International Journal of Digital Society (IJDS), Volume 1*, 2010.

[21] B. Auernheimer and M. Tsai, "Biometric authentication for web-based course examinations," in *Proc. of the 38th Int. Conf. on System Sciences*, 2005.

[22] M. Ramim and Y. Levy, "Towards a framework of biometrics exam authentication in e-learning environments," in *Proc. of the Information Resources Management Association Conf. (IRMA)*, 2007, pp. 539–543.

[23] O. Pieczul and S. Foley, "Collaborating as normal: detecting systemic anomalies in your partner," in *Proc. of 22nd International Security Protocols Workshop, Cambridge*, 2014, (to appear).

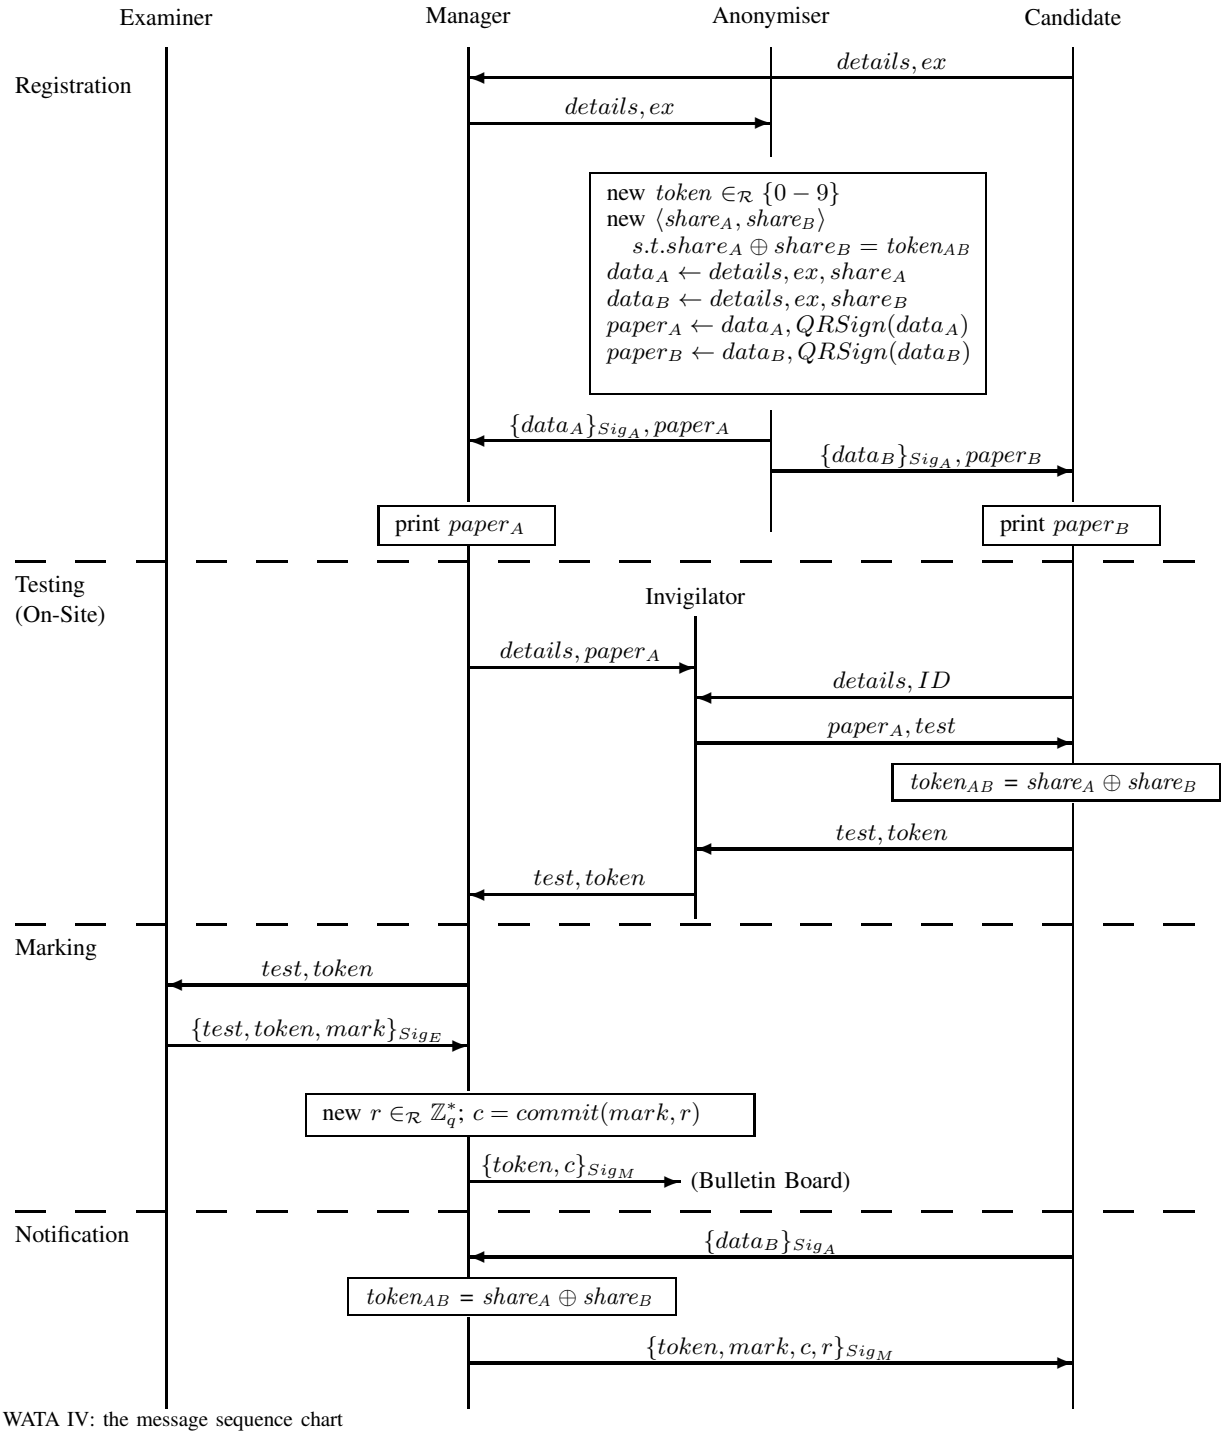

Fig. 2. WATA IV: the message sequence chart