MAGICARPP: ModulAr GeneralIzed Check of Age Requirement while Preserving Privacy

Martin Bieri[1], Olivier Blazy[2], Solenn Brunet[1], and Jerome Gorin[3]

[1]LINC, CNIL, Paris, France *
[2]École Polytechnique, Palaiseau, France
[3]UniLasalle, Amiens, France

Abstract

Age verification is increasingly mandated by law, especially in online services. In this paper, we present a proof-of-concept protocol for privacy-preserving age verification. Our approach allows users to prove compliance with age requirements without revealing their identity to the service provider. Furthermore, the website remains unaware of which authority issued the age certification, and the certification authority does not learn which website is making the request or the specific age being verified. This design ensures strong privacy guarantees for all parties involved. Its properties served as a baseline for part of the reference document published by the French regulator ARCOM.

1 Introduction

Online age verification consists of ensuring that individuals browsing a site or viewing a content meet the minimum age threshold for its access. Generally speaking, this means proving that you are over the age of majority. However, online age verification systems may concern a wider variety of use cases for which the verification is not linked to majority. For instance, age limits apply to different age thresholds (13, 15, 16, etc.) for social networks and video games.

In this paper, we present a generic framework to provide *double anonymity*. This concept was made into the reference documents [1], to regulate age-based access control to adult websites in France.

This paper describes privacy technics that could conceal the true identity of individuals to a website during an age verification process. Specifically, these technics ensure that the true identity of a user is hidden from the publishers of online services, who must put in place mechanisms to prevent underage from accessing content. Instead, the age verification process is delegated to ad-hoc Identity Providers (IdPs), which is trusted by publishers for implementing state-of-the-art age verification technics. If the user has the required age, a pseudonymous proof-of-works is generated and transmitted to the publisher. The service provider does not know which IdPs certified the age of the user, and the IdP does not know which website is requesting an age check, and which age is requested.

1.1 The legal context

The growing number of studies and discussions concerning online age verification is part of a wider issue: the protection of children online [29].

This topic has been recently receiving renewed attention, especially in Europe [15, 9]. One of the reason of this increasing concerns is the easy access to on-line pornographic content, in particular, with the growing importance of the "Tubes", those large sites whose content is, for the most part at any rate, freely available [21]; and the growing proportion of underage consulting them on a regular basis and at an earlier age [30].

At the European level, a legal framework is already existing that can apply to these systems of age verification. For instance, the GDPR [12] applies to every processing of personal data. In this framework, two articles are directly linked to age verification systems: article 8 (age of digital consent) or article 9 (sensitive data). More recently, the Digital Services Act [13] mentions age verification in article 35. Article 28 also states that online platforms must put in place measures to protect children and the European Commission published in July 2025 guidelines to protect children online [7], in which it recommends the use of effective age assurance methods. On the broader front, the European "Better Internet for Kids" strategy is a good summary of the issues that needs to be tackled [8].

*The views and opinions expressed in this paper do not necessarily express the views of the CNIL or any individual commissioner

There are also national obligations towards the publishers of these sites, who must put in place mechanisms to prevent underage from accessing content that is forbidden to them. In practice, the mechanisms put in place by sites are usually a simple self-declaration of majority, often by clicking on an "I'm over 18" button or entering a suitable date of birth, which is easily circumventable (and circumvented).

The UK is one of the first countries to put back this issue in 2017, with the "Digital Economy Act" bill [22]. The third part of this bill includes obligation for pornographic websites to implement alternative online age verification systems. These obligations not only sparked debate, but also generated a number of initiatives by third-party players offering these other systems, such as cards to be purchased in person, providing an identifier and a code. Provisions were initially postponed, and were recently implemented, the "Online Safety Bill" [4].

In France, a law [20] enacted in 2020 specifies guidelines for online age verification in the case of accessing pornographic contents: self-declaration mechanisms are explicitly forbidden. It is therefore up to the site publisher to find a new system, or risk having his or her website blocked - the supervisory role then falls to the Autorité de Régulation de la Communication Audiovisuelle [1] et Numérique (Arcom, formerly Conseil Supérieur de l'Audiovisuel, CSA), which can then ask a judge to order blocking at a later date.

1.2 Motivations

Age verification, by the very nature of the Internet, is confronted with the difficulty of knowing "who is behind" the terminal being used with a certain level of confidence

Yet, an age verification processes often rely on an identity document and a checking-process that ensures the person presenting this document actually corresponds to the identity presented.

However, linking an individual identity to their online activity raises many privacy risks, especially in the case where the accessed content could be linked to sensitive data such as natural person's sex life or sexual orientation. Unlike many other online uses where the true identity is required for legal reasons (e.g. online purchases, gambling), age verification does not necessarily imply this prior identification.

Our privacy-preserving system for age verification systems therefore meet the following criteria: system reliability –to ensure that as few underage individuals are allowed and that the system cannot be easily bypassed– , and security –in terms of protecting all individuals and their personal data–.

As such this privacy-preserving solution would allow to guarantee users' right while satisfying legislators expectations. This also touches on the issue of freedom of information – and, more generally, runs up against the question of the Internet's open model, freely accessible to users and site publishers alike; and limits the issues of self-censorship or "chilling effect" as the solution is auditable, and user are free to choose the age verifier they trust instead of being tunneled to a single one.

It seems out of scope to redesign the Internet protocol, or enforce usage of protocols like Tor, as such protecting users against their internet access providers seems an unreasonable goal. Besides this, we want to offer as much protection as possible, in particular the content providing website should not learn anything new from the protocol (if the users accesses the website via classical browsing it will leak its IP address, but we aim at not giving away its name, its real age or even the age-check authority he used), in the same way, we also want to protect the user by hiding from the age authority which website is requesting an age check, or even which age is being checked.

As such, we can not claim to be "anonymous" [2], as IP address is an identifying information in the eyes of the law, but we are as pseudonymous as possible as long as we do not alter the routing protocol.

1.3 Related Work

This work has already been used as a proof-of-concept (PoC) [24]. The principle of the PoC has been notably taken up by the government, which in February 2023 called on industry and digital players to seize upon it to experiment with solutions based on this mechanism [27].

Delegating the access of users to a service using a third-party permission is relatively well explored topic, and many protocols are already designed for handling this process securely.

For instance, the OAuth 2.0 protocol [19], that replaces and obsoletes OAuth 1.0, has been originally designed for applications to access a user's data securely, without requiring the user handing over an account password. Instead, identity attributes are stored by identity providers (IdPs) and provided to relying party (RP) websites as required. In the same way, the protocol OpenID Connect [28] is built on top of OAuth 2.0 to manage federated authentication.

Both protocols have been designed with security in mind, but with limited consideration for privacy. For example, [23] explains that the IdP could collect many privacy-related information from the RP, based on the **redirect_uri** or the **client_id**. In the same a way, [16] explains that redirection in OAuth could reveal the identities of the websites visited by users to the IdP, including metadata such as time of visits. More generally, both protocols requires that

[1] The french authority for audiovisual and digital communication

[2] "Double-Anonymity" was coined by policy makers but is not perfectly GDPR appropriate.

applications register with the authorization server so that API requests are able to be properly identified, which is not a desired in our use case.

More recently, the Privacy Pass [10] protocol has been defined to allow some proof-of-works to be used for authenticating to service, while retaining anonymity for the user via the use of a cryptographic protocol.

Although this solution can be viewed as a good candidate for generating age requirement proof, this protocol allows only a single, and known by the service, IdP to be involved in the signature process. Conversely, our contribution make use of Group Signatures which allows multiple IdPs to be transparently used for a given age verification process. This feature guarantees that no essential IdPs are in placed for accessing RPs, which could be aware of some habits of its users, as well as enabling revocation of any these IdPs by an opener if required.

2 Preliminaries

In this section, we give notations and cryptographic primitives that we use in this paper. This is the occasion to stress inputs and outputs properties of the algorithm, and detail the required security properties.

2.1 Cryptographic Primitives

2.1.1 Group Signatures (GS)

A group signature scheme [6] is a protocol which lets a member of a group individually issue signatures on behalf of the group, in an anonymous but revocable way: an opener is able to revoke anonymity of the actual signer in case of abuse. The members of the group are not trusted. Bellare, Shi and Zhang [2] extended the initial model to dynamic groups (the BSZ model), emphasizing the importance of unforgeability and anonymity.

In addition to *unforgeability*, Group signatures guarantee *anonymity*, which means that nobody (except the opener) can link the signature to the signer, but also *unlinkability*, which means that one cannot tell whether two signatures have been produced by the same user.

We use similar notations as [2] for the BSZ model to define, in a game-based way, the security notions. In a group signature scheme, there are several users, which are all registered in a PKI. We thus assume that each user $\mathcal{U}_i$ owns a pair $(\mathsf{usk}[i], \mathsf{upk}[i])$ certified by the PKI. There is a group manager, also known as *Issuer*, since he will issue certificates to grant access to the group, and an *Opener* that will be able to revoke anonymity, and thus trace back the actual signers. Those two authorities are not necessarily the same. To be precise, a group signature scheme is a sequence of (interactive) protocols:

Defitinion 2.1 (Group Signature) $\mathsf{GS} = (\mathsf{Setup}, \mathsf{Join}, \mathsf{Sign}, \mathsf{Verif}, \mathsf{Open}, \mathsf{Judge})$*:*

- $\mathsf{Setup}(1^{\mathfrak{K}})$*: this algorithm generates the global parameters of the system, the public key* pk *and the private keys: the master secret key* msk *given to the group manager, and the opening key* skO *sent to the opener;*
- $\mathsf{Join}(\langle \mathcal{U}_i(\mathsf{usk}[i]), \mathcal{M}(\mathsf{msk})\rangle)$*: this is an interactive protocol between a user* $\mathcal{U}_i$ *(using his secret key* $\mathsf{usk}[i]$*) and the group manager* $\mathcal{M}$ *(using his private key* msk*). At the end of the protocol, the user obtains a signing key* $\mathsf{sk}[i]$ *(or group membership certificate), and the group manager adds the user to the registration list, storing some information in* $Reg[i]$
- $\mathsf{Sign}(\mathsf{sk}[i], m; \mu)$*: To sign a message* m*, the user uses his secret key* $\mathsf{sk}[i]$ *and some randomness* μ*, to output a signature* σ *(valid under the group public key* pk*)*
- $\mathsf{Verif}(\mathsf{pk}, m, \sigma)$*: anybody should be able to verify the validity of the signature* σ *on the message* m*, w.r.t. the public key* pk*. This algorithm thus outputs 1 if the signature is valid, and 0 otherwise*
- $\mathsf{Open}(\mathsf{skO}, \mathsf{pk}, m, \sigma)$*: granted the opening key* skO*, for a valid signature* σ *w.r.t. the public key* pk*, the* Opener *can provide the identity signer. It thus outputs the user* i*, together with a proof* Π
- $\mathsf{Judge}(\mathsf{pk}, m, \sigma, i, \Pi)$*: this algorithm publicly checks the claim of the opener*

2.1.2 Zero-Knowledge Proofs (ZK)

Zero-Knowledge Proofs are a powerful tool introduced by Goldwasser, Micali and Rackoff in [17]. They let a user prove the veracity of a statement $\mathcal{S}$ without leaking any additional information. They have found many applications in various protocols since (Anonymous Credentials, Anonymous Signatures (Group Signatures, Ring Signatures, Blind Signatures,...), Online Voting, PAKE, Proof of a shuffle, ...).

Such proofs are expected to have three properties:

- **Completeness:** If $\mathcal{S}$ is true, the honest verifier will be convinced of this fact.

- **Soundness:** If $\mathcal{S}$ is false, no cheating prover can convince the honest verifier that it is true except with negligible probability.
- **Zero-knowledge:** Anything that is feasibly computable from the proof is also feasibly computable from the assertion itself .

An upgrade to this kind of proofs can be obtained by removing the interaction between the prover and the verifier. In this case we speak about NIZK, Non-interactive Zero-Knowledge Proof. Fiat Shamir [14] presented an heuristic showing how to transform Interactive proofs into Non-Interactive ones. Several approaches have been proposed since ([3],[11],...), but they are all rather inefficient, until Groth-Sahai methodology in [18] to prove pairing equations.

2.1.3 Blind Signature (BS)

Signatures can be extended in several ways. One of them is quite useful in electronic votes, and e-cash protocols, when someone might want an authority (Bank, Poll centre, ...) to sign a specific message he wants to keep secret. For that, we use the notion of *Blind Signature* introduced by Chaum [5] for electronic cash in order to prevent the bank from linking a coin to its spender.

Such protocol can easily be derived from digital signatures. Instead of having a signing phase $\mathsf{Sign}(\mathsf{sk}, M; \mu)$ we have an interactive phase $\mathsf{BSProtocol}\langle\mathcal{S},\mathcal{U}\rangle$ between the user $\mathcal{U}(\mathsf{vk}, M; \rho)$ who will transmit a masked challenge under some randomness ρ in order to obtain a signature valid under the verification key vk, and the signer $\mathcal{S}(\mathsf{sk}; \mu)$, who will generate something based on this value, and his secret key which should lead the user to a valid signature.

Such signatures are correct if when both the user and signer are honest then $\mathsf{BSProtocol}\langle\mathcal{S},\mathcal{U}\rangle$ does indeed lead to valid signature on M under vk.

There are two additional security property, one protecting the signer, the other the user:

- On one hand, there is an *Unforgeability* property, where a malicious user shouldn't be able to compute $n+1$ valid signatures on different messages after at most n interactions with the signer.
- On the other hand, the *Blindness* property says that a malicious signer who signed two messages M_0 and M_1 shouldn't be able to decide which one was signed first.

3 Our proposed framework

Our primitive will consider 4 kinds of participants. A master authority M, that will enable age-checking authorities A. This authority M should be able to interact with age-checking entities, and provide each of them with a different group signing key $\mathsf{usk}[A]$.

Users U through various means will register with the authorities A (opening a bank account, get a nation-wide digital id, ...). And when sending a challenge to an authority, they expect to get a certified answer stating whether they are above an age requirement.

Content Websites (C), interacts with users U, and send them challenges to learn whether they are allowed to access to a specific content.

As this protocol is targeted for an internet deployment, we assume the existence of an independent PKI, ensuring to perfectly authenticate and send private messages to all the authorities (M and A), this is classically achieved through SSL/TLS but the specifics are not a requirement for the protocol. For simplification, we assume everywhere the existence of authenticated secure channel.

This setup lead us to a first version of the protocol indicating the global workflow, and the main technique behind the "double anonymity".

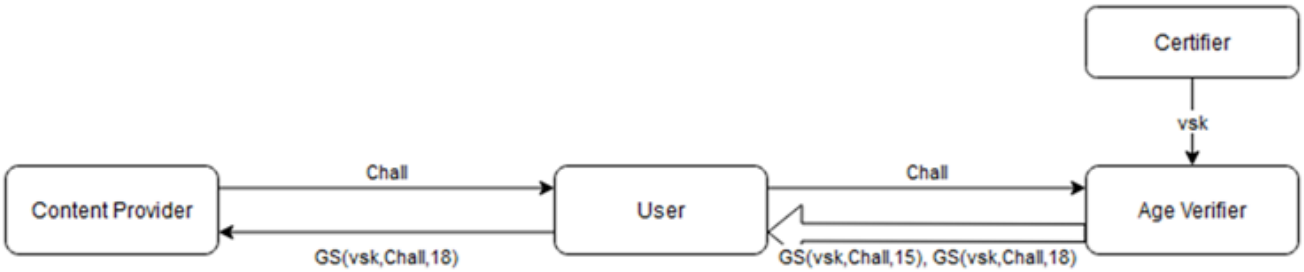

Figure 1: Workflow of a naive request

In this context, when trying to access content, the user receives a challenge, forwards it to a certified Age Verifier of its choosing. He then receives group signatures on the challenge for the various age threshold he respects.

The User then picks the threshold adapted for the content he is trying to access, and sends it to the content provider that can check that the challenge was signed for the given threshold by a certified age verifier.

It should be noted, that while this design achieves part of the expected properties it still fails at some point:

- The challenge creates a *pseudonym* that may allow colluding content provider and age verifier to track a user habit
- Age providers lack a capacity to track and monetize age verification directly to content providers
- Users are not guaranteed that tracking data is not present

As such, we needed to add extra mechanisms in the design

- Blind Signatures ensures that the age verifier does not learn the challenge
- Randomizable signatures ensure the user can remove potential tracking data
- A threshold batch opening allows to count the number of verifications provided by an age verifier in a given timeframe

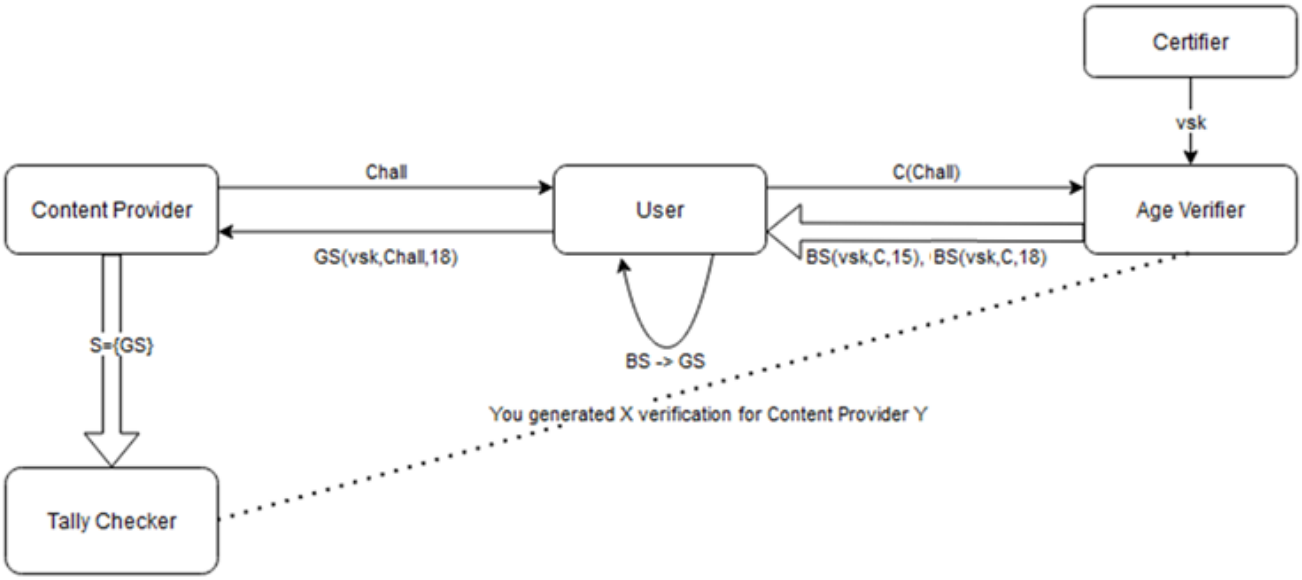

Figure 2: Workflow of the advanced scheme

The workflow is vaguely the same. The user received a challenge. He then commits to the challenge, sends it to the certified age verifier, that does a blind signature on it. The user picks the threshold he wants, unblind the signature, and transforms it into a group signature, that he then sends to the content provider, that can verifies it.

At various interval, content provider can then send the aggregated received tokens to an authority, that can do a threshold opening to count how many of those were provided by a given age verifier. (Adding canari tokens can ensure content providers don't hide part of their tokens to pay less...). This feature was detailed in [25].

4 Concluding remarks

Implementation We implemented our primitive in C using the PBC (Pairing-Based Cryptography) library [26]. We provide some wrappers around this library in Python (etc.). Each entity has access to python wrapper / API to have a plug and play access to the protocol.

The naive protocol code is available here `https://gitlab.adullact.net/linc/siggroup`, and is available under a GNU GPL V3.0 license. The running time of each algorithm is negligible (41ms for signing, 28ms for verification on a classical laptop: Processor Intel(R) Xeon(R) E-2286M CPU @ 2.40GHz) making sure that it can be integrated in the user experience without disrupting their navigation.

The primitives chosen have many known instantiations, ensuring some flexibility in the proposed framework. In particular, there exist post-quantum variants, allowing if required a quantum-resilient pseudonymity.

Context The solution presented in the paper served as a basis for the French guidelines published by the Arcom [1] for age verification. It provides a proof of concept showing that a high level of security could be reach, without relying on secure elements on a device (and so without requiring users to have their smartphone at hand, while accessing restricted websites).

Assuming, a readily available trusted environment some interactions can be removed (a local TEE/TPM could generate the signature on the challenge), however the design ensure the user does not need to trust the secure computation. It also shows how group signatures can be leveraged to give the user a choice in which age verification they use, without leaking this information to the content provider.

Enabling the user to have a choice will reduce defiance against such policy at no extra cost.

The API is designed in such a way, that it can be compatible with several forms of age verification, including (but not limited) the future European digital identity wallet.

Further Evolutions Of course, this is just a proof of concept.

Adding, the blind signature allowed further privacy enhancement, and also permits the user to check whether the age verifier is behaving correctly, of if they signed under an alternative key. We also added a revocation mechanism that, without breaking the anonymity of the age verifier used, allows checking whether it is one that has been forbidden (because of misbehavior).

The group signature used, allows to create a hierarchy where a national agency would certify age verifiers allowed to operate in a country. This could be further refined, with age verifier for given thresholds, or the hierarchy could be done at a continent / worldwide level, where an entity would certify national agencies in charge of certifying residents. This could pave the way to a global age verification service making the use of VPNs moot for the purpose of avoiding age verification while protecting the user real country.

References

[1] ARCOM. Référentiel technique sur la vérification de l'age pour la protection des mineurs contre la pornographie en ligne. Arcom website (In French), 2024.

[2] Mihir Bellare, Haixia Shi, and Chong Zhang. Foundations of group signatures: The case of dynamic groups. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 136–153. Springer, Heidelberg, February 2005.

[3] Manuel Blum, Paul Feldman, and Silvio Micali. Proving security against chosen cyphertext attacks. In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 256–268. Springer, Heidelberg, August 1990.

[4] Anthony Burton, Mark Soames, and Alexandra Cohen. The online safety bill 2022. *Solic. J.*, 165:58, 2022.

[5] David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 199–203. Plenum Press, New York, USA, 1982.

[6] David Chaum and Eugène van Heyst. Group signatures. In Donald W. Davies, editor, *EUROCRYPT'91*, volume 547 of *LNCS*, pages 257–265. Springer, Heidelberg, April 1991.

[7] European Commission. Guidelines on measures to ensure a high level of privacy, safety and security for minors online. Online, 2025.

[8] European Commission, Content Directorate-General for Communications Networks, and Technology. *The European strategy for a better internet for kids (BIK+)*. Publications Office of the European Union, 2022.

[9] Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee, and the Committee of the Regions. A digital decade for children and youth: the new european strategy for a better internet for kids (bik+). `https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52022DC0212`, 2022.

[10] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *Proc. Priv. Enhancing Technol.*, 2018(3):164–180, 2018.

[11] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 566–598. Springer, Heidelberg, August 2001.

[12] European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council.

[13] European Parliament and Council of the European Union. Regulation (EU) 2022/2065 of the European Parliament and of the Council.

[14] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.

[15] EU Funding. Outline and trial an infrastructure dedicated to the implementation of child rights and protection mechanisms in the online domain. `https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/pppa-agever-01-2020`, 2020.

[16] Srivathsan Morkonda Gnanasekaran. *USER PRIVACY IN OAUTH-BASED SINGLE SIGN-ON SYSTEMS*. PhD thesis, CARLETON UNIVERSITY Ottawa, 2024.

[17] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

[18] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 415–432. Springer, Heidelberg, April 2008.

[19] Dick Hardt. The oauth 2.0 authorization framework. Technical report, 2012.

[20] JORF. Art. 23 de la loi n° 2020-936 du 30 juillet 2020 visant à protéger les victimes de violences conjugales. JUSX1935275L, 2020.

[21] Patrick Keilty. Desire by design: pornography as technology industry. *Porn Studies*, 5(3):338–342, 2018.

[22] Legislation.go.uk. Digital economy act. `https://www.legislation.gov.uk/ukpga/2017/30/contents/enacted`, 2017.

[23] Wanpeng Li and Chris J Mitchell. User access privacy in oauth 2.0 and openid connect. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 664–6732. IEEE, 2020.

[24] LINC. Demonstration of a privacy-preserving age verification process. `https://linc.cnil.fr/demonstration-privacy-preserving-age-verification-process`, 2022.

[25] LINC. [follow-up] - age verification: the economic argument. `https://linc.cnil.fr/en/follow-age-verification-economic-argument`, 2023.

[26] Ben Lynn. Pbc library-pairing-based cryptography. *http://crypto. stanford. edu/pbc/*, 2007.

[27] Vie publique. Déclaration de m. jean-noël barrot, ministre chargé de la transition numérique et des télécommunications, sur la protection des mineurs en ligne. `https://www.vie-publique.fr/discours/288266-jean-noel-barrot-14022023-protection-des-mineurs-en-ligne`, 2023. [Accessed 24-04-2024].

[28] Natsuhiko Sakimura, John Bradley, Mike Jones, Breno De Medeiros, and Chuck Mortimore. Openid connect core 1.0. *The OpenID Foundation*, page S3, 2014.

[29] David Smahel, Hana Machackova, Giovanna Mascheroni, Lenka Dedkova, Elisabeth Staksrud, Kjartan Ólafsson, Sonia Livingstone, and Uwe Hasebrink. Eu kids online 2020: Survey results from 19 countries, 2020.

[30] Gila Cohen Zilka. Awareness of esafety and potential online dangers among children and teenagers. *Journal of Information Technology Education. Research*, 16:319, 2017.