The Apple PSI Protocol

Mihir Bellare
Department of Computer Science and Engineering
University of California, San Diego

July 30, 2021

Writer bio: Mihir Bellare is a cryptographer whose work has concerned proof-based assurance for practical cryptography. He obtained his BS at Caltech (1986) and his PhD at MIT (1991) and has 30 years of experience in cryptography. He is a Fellow of the ACM and IACR, and a recipient of the ACM Paris Kanellakis Theory and Practice Award and the Levchin prize. He co-introduced the random oracle model in cryptography. Cryptographic algorithms that he has co-developed, and are widely used in practice, include HMAC, RSA-OAEP, RSA-PSS and DHIES.

The National Center for Missing and Exploited Children (NCMEC) explains that "United States federal law defines child pornography as any visual depiction of sexually explicit conduct involving a minor [1]." They refer to these images as Child Sexual Abuse Material (CSAM). They document the harm they cause and note that "the disturbing reality is that the Internet platforms we use every day ... are now being used to ... collect CSAM."

Apple is aiming to limit CSAM on its platforms. Apple users (also called clients) store photos in iCloud. Apple would like to detect if any of these photos belongs to NCMEC's database of CSAM photos. If the number of these matches exceeds some pre-determined threshold, indicating systematic presence of CSAM, Apple will report the user to appropriate authorities.

Taking action to limit CSAM is a laudable step. But its implementation needs some care. Naively done, it requires scanning the photos of all iCloud users. But our photos are personal, recording events, moments and people in our lives. Users expect and desire that these remain private from Apple. Reciprocally, the database of CSAM photos should not be made public or become known to the user.

Apple has found a way to detect and report CSAM offenders while respecting these privacy constraints. When the number of user photos that are in the CSAM database exceeds the threshold, the system is able to detect and report this. Yet a user photo that is not in the CSAM database remains invisible to the system, and users do not learn the contents of the CSAM database.

This is done using cryptography. Apple starts from a well-established cryptographic tool called Private Set Intersection (PSI). It then enhances it, to add further privacy (PSI would already provide Apple only with the user CSAM photos in the database, but the Apple protocol further denies it even knowledge of matches when the number of them is below threshold) and to satisfy some system and performance constraints.

For the protocol to provide the desired privacy, the cryptography needs to be right. How do we know that the math works, meaning that the privacy goals are met? Cryptographers assess this by giving what are called proofs of security. Such a proof establishes that the protocol meets a certain mathematical *definition* of security, assuming building blocks used within (these include AES and elliptic-curve cryptography) are themselves secure.

There is enough conviction in the community that proofs are important that they tend to be a requirement for protocols to be standardized. Apple has, accordingly, sought such proofs for its protocol. Their document [3] gives one such proof.

I complement this, in a companion document [2], with another proof. It uses different proof methods, and, most importantly, gives what cryptographers call a concrete-security analysis. This evaluates security quantitatively, giving evidence that the protocol is not only secure in principle, but is so for the key sizes in actual use.

Why another analysis? The Apple protocol will see hundreds of millions of uses. It is desirable that it receive extensive analysis, done by different people, using different methods.

References

[1] National Center for Missing and Exploited Children (NCMEC), June 2021. https://www.missingkids.org/theissues/csam.

[2] M. Bellare. A concrete security analysis of the Apple PSI protocol, July 2021.

[3] A. Bhowmick, D. Boneh, S. Myers, K. Talwar, and K. Tarbe. The Apple PSI system, July 2021.