Proofs of Storage: Theory, Constructions and Applications

Seny Kamara

Microsoft Research
senyk@microsoft.com

Abstract. Proofs of storage (PoS) are cryptographic protocols that allow a client to efficiently verify the integrity of remotely stored data. To use a PoS, the client sends an encoded version of its data to the server while keeping a small amount of state locally. At any point in time, the client can then verify the integrity of its data by executing a highly-efficient challenge-response protocol with the server.

Since their introduction in 2007 by Ateniese et al. (Computer and Communications Security, 2007) and Juels and Kaliski (Computer and Communications Security, 2007), PoS have received a lot of attention from the research community. This is due in large part to their potential practical applications (e.g., to the design of various kinds of secure cloud storage systems) but also due to their inherent theoretical properties and their connections to fundamental primitives like digital signatures, identification schemes, zero-knowledge proofs and error-correcting codes.

In this talk, I will survey the current state of PoS research. This will include the many variants of PoS that have been invented over the years, how to design them, the connections that have been established between PoS and other primitives and the many new applications PoS have enabled.

References

1. Ateniese, G., Burns, R., Curtmola, R., Herring, J., Kissner, L., Peterson, Z., Song, D.: Provable data possession at untrusted stores. In: ACM Conference on Computer and Communication Security (CCS 2007). ACM (2007)
2. Ateniese, G., Di Pietro, R., Mancini, L., Tsudik, G.: Scalable and efficient provable data possession. In: Conference on Security and Privacy in Communication Networks (SecureComm 2008), pp. 9:1–9:10 (2008)
3. Ateniese, G., Faonio, A., Kamara, S., Katz, J.: How to authenticate from a fully compromised system (under submission, 2013)
4. Ateniese, G., Kamara, S., Katz, J.: Proofs of storage from homomorphic identification protocols. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 319–333. Springer, Heidelberg (2009)
5. Benson, K., Dowsley, R., Shacham, H.: Do you know where your cloud files are? In: ACM Cloud Computing Security Workshop (CCSW 2011), pp. 73–82 (2011)
6. Bowers, K., Juels, A., Oprea, A.: Proofs of retrievability: Theory and implementation. Technical Report 2008/175. Cryptology ePrint Archive (2008)

T. Muntean, D. Poulakis, and R. Rolland (Eds.): CAI 2013, LNCS 8080, pp. 7–8, 2013.

7. Bowers, K., Juels, A., Oprea, A.: HAIL: a high-availability and integrity layer for cloud storage. In: ACM Conference on Computer and Communications Security (CCS 2009), pp. 187–198. ACM (2009)
8. Bowers, K., van Dijk, M., Juels, A., Oprea, A., Rivest, R.: How to tell if your cloud files are vulnerable to drive crashes. In: ACM Conference on Computer and Communications Security (CCS 2011), pp. 501–514. ACM (2011)
9. Cash, D., Küpçü, A., Wichs, D.: Dynamic proofs of retrievability via oblivious RAM. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 279–295. Springer, Heidelberg (2013)
10. Dodis, Y., Vadhan, S., Wichs, D.: Proofs of retrievability via hardness amplification. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 109–127. Springer, Heidelberg (2009)
11. Erway, C., Küpçü, A., Papamanthou, C., Tamassia, R.: Dynamic provable data possession. In: ACM Conference on Computer and Communications Security (CCS 2009), pp. 213–222. ACM, New York (2009)
12. Gondree, M., Peterson, Z.: Geolocation of data in the cloud. In: ACM Conference on Data and Application Security and Privacy (CODASPY 2013), pp. 25–36. ACM (2013)
13. Juels, A., Kaliski, B.: PORs: Proofs of retrievability for large files. In: Ning, P., De Capitani di Vimercati, S., Syverson, P. (eds.) ACM Conference on Computer and Communication Security (CCS 2007). ACM (2007)
14. Juels, A., Oprea, A.: New approaches to security and availability for cloud data. Communications of the ACM 56(2), 64–73 (2013)
15. Naor, M., Rothblum, G.: The complexity of online memory checking. In: IEEE Symposium on Foundations of Computer Science (FOCS 2005), pp. 573–584. IEEE Computer Society (2005)
16. Naor, M., Rothblum, G.: The complexity of online memory checking. Journal of the ACM 56(1), 2:1–2:46 (2009)
17. Peterson, A., Gondree, M., Beverly, R.: A position paper on data sovereignty: The importance of geolocating data in the cloud. In: USENIX Workshop on Hot Topics in Cloud Computing, HotCloud 2011 (2011)
18. Shacham, H., Waters, B.: Compact proofs of retrievability. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 90–107. Springer, Heidelberg (2008)
19. Stefanov, E., van Dijk, M., Juels, A., Oprea, A.: Iris: a scalable cloud file system with efficient integrity checks. In: Annual Computer Security Applications Conference (ACSAC 2012), pp. 229–238. ACM, New York (2012)
20. van Dijk, M., Juels, A., Oprea, A., Rivest, R., Stefanov, E., Triandopoulos, N.: Hourglass schemes: how to prove that cloud files are encrypted. In: ACM Conference on Computer and Communications Security (CCS 2012), pp. 265–280. ACM (2012)
21. Wang, C., Chow, S., Wang, Q., Ren, K., Lou, W.: Privacy-preserving public auditing for secure cloud storage. Cryptology ePrint Archive, Report 2009/579 (2009), `http://eprint.iacr.org/2009/579`
22. Watson, G., Safavi-Naini, R., Alimomeni, M., Locasto, M., Narayan, S.: LoSt: location based storage. In: ACM Cloud Computing Security Workshop (CCSW 2012), pp. 59–70. ACM (2012)