

Accelerating PIR


Access Privacy on Untrusted Storage

Peter Williams, Miroslava Sotakova, Radu Sion.

petertw@cs.stonybrook.edu, gwhitehawk@gmail.com, sion@cs.stonybrook.edu


The Scenario

Storage is most cost-effectively maintained in bulk by dedicated providers.




But why should we trust the provider, who may have variable and adverse financial incentives, or legal obligations?

So we encrypt our data before sending it to the provider.



Encryption is not enough, however!




The storage provider sees:

```
07:35 AM. Uploaded encrypted document 734.
07:35 AM. Updated encrypted keyword indexes #'s 102, 592, 1002.
11:30 AM. Company announces merger
11:35 AM. Many accesses to keyword index # 592.
-> Does document 7 contain merger information?
```



The untrusted storage provider can link access patterns to external knowledge!

A Solution

Main idea from [1]: client periodically reshuffles the database.




Our contribution: a small amount of client storage leads to big gains in reshuffle performance. Perform a merge sort, but buffer the reads at the client to hide the permutation. The read cursors remain close, since keys are uniformly random.



References:

- [1] Oded Goldreich and Rafail Ostrovsky. Software Protection and Simulation on Oblivious RAM. *Journal of the ACM*, 45:431-473, May 1996.
- [2] A. Iliev and S.W. Smith. Private Information Storage with Logarithmic-space Secure Hardware. In *Proceedings of i-NetSec 04: 3rd Working Conference on Privacy and Anonymity in Networked and Distributed Systems*, pages 201-216, 2004.
- [3] Shuhon Wang, Xuhua Ding, Robert H. Deng, and Feng Bao. Private Information Retrieval Using Trusted Hardware. In *Proceedings of the European Symposium on Research in Computer Security ESORICS*, pages 49-64, 2006.
- [4] Peter Williams, Radu Sion. Usable PIR. *Proceedings of the 2008 Network and Distributed System Security (NDSS) Symposium*, 2008.

This notion of regularity through uniformity of randomness allows our small client buffer to improve performance throughout the entire reshuffle process:



Amortized Query Complexity:

For client storage $O(\sqrt{n} \log n)$

- Goldreich/Ostrovsky [1] - $O(\log^4 n)$
- Smith/Illiev [2] - $O(\sqrt{n} \log n)$
- Wang et al. [3] - $O(\sqrt{n} / \log n)$
- This protocol [4] - $O(\log^2 n)$

Predicted performance:

