

Revisiting Square Root ORAM

Efficient Random Access in Multi-Party Computation



Samee Zahur
Jack Doerner
David Evans



Xiao Wang
Jonathan Katz

Yale


Mariana Raykova



Adrià Gascón

oblivc.org/sqoram


Secure multi-party computation applications




Set intersection
[FNP04]




Iris code matching
[LCPLB12]



Matrix factorization for
recommendations
[NIWJTB13]




Median computation
[AMP04]



Linear ridge-regression
[NWIJBT13]

Random Access




```
void oscrypt_smix(obliv uint8_t * B, s
...
for (i = 0; i < N; i += 2) {
    j = integerify(X, r) & (N - 1);
    temp = V2[j];
    xorBits(X, temp, 32*r);
    oscrypt_blockmix_salsa8(X, Y, r);

    j = integerify(Y, r) & (N - 1);
    temp = V2[j];

    for (size_t jj = 0; jj < 32 * r; j
        Y[jj] ^= temp[jj];
    }
```


Hiding access pattern

Linear scan



Access every element
Per-access cost: $\Theta(n)$

Oblivious RAM



Continually shuffle elements around
Per-access cost: $\Theta(\log^p n)$

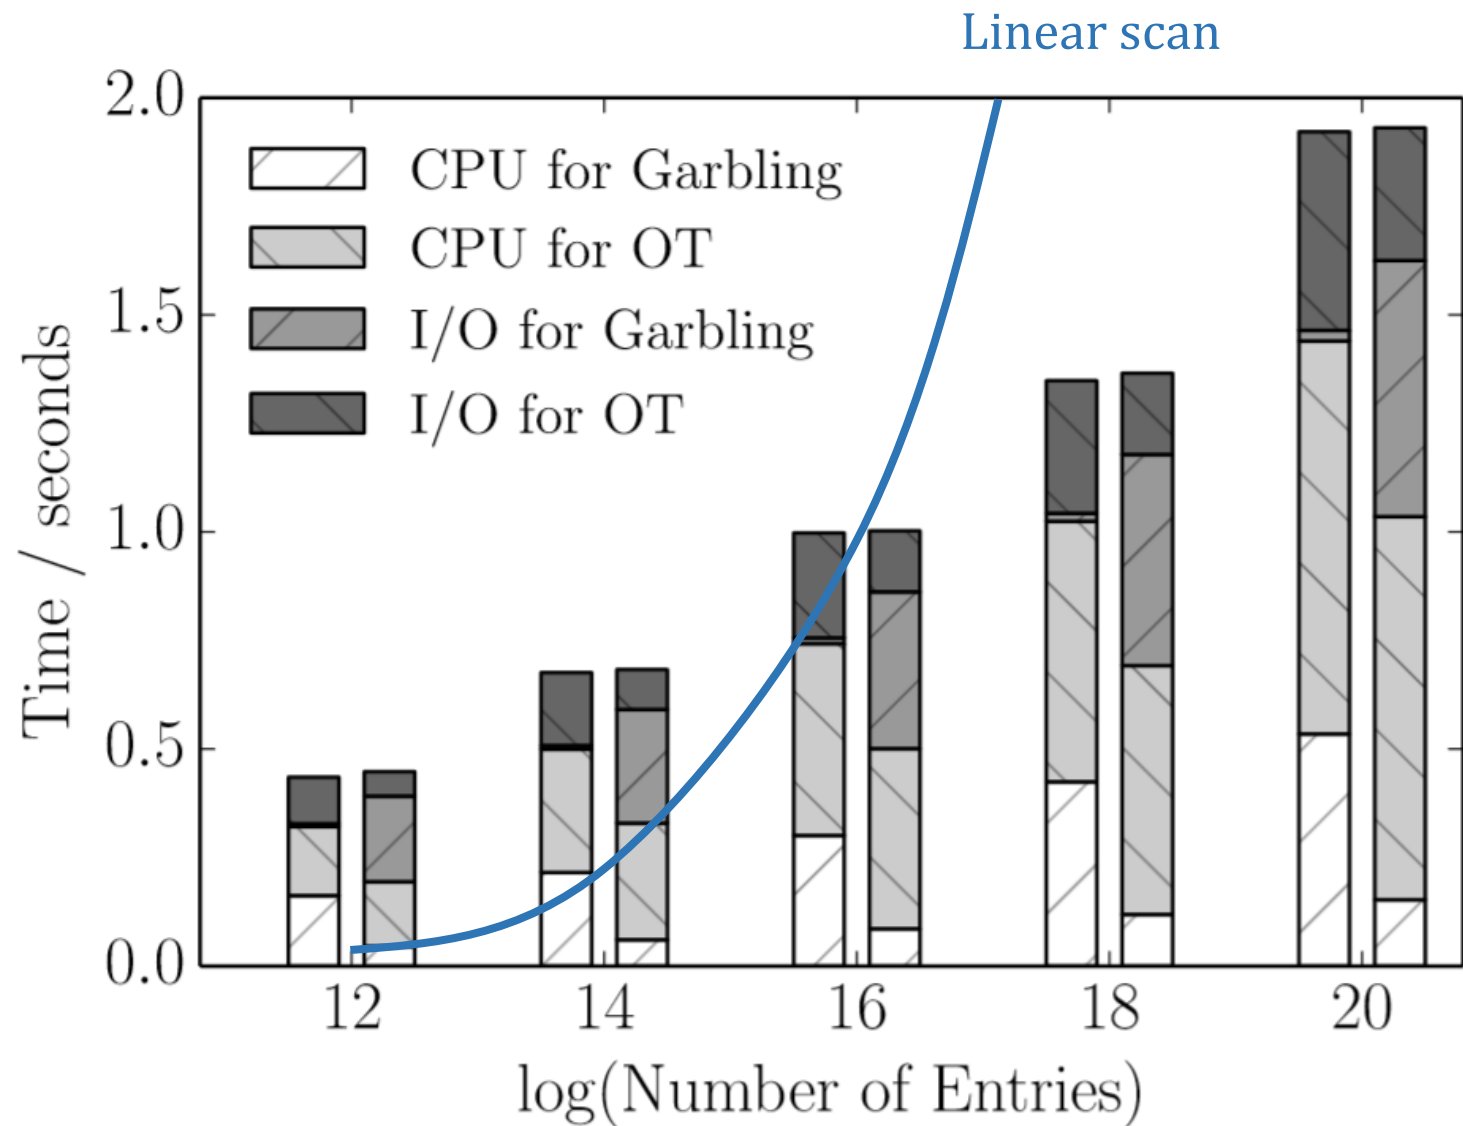



Figure from: Wang, Chan, Shi. **Circuit Oram**. CCS'15




Approach: revisit old schemes


Classic “square root” scheme by Goldreich and Ostrovsky (1996).

Considered slow for MPC because of per-access hash evaluation.

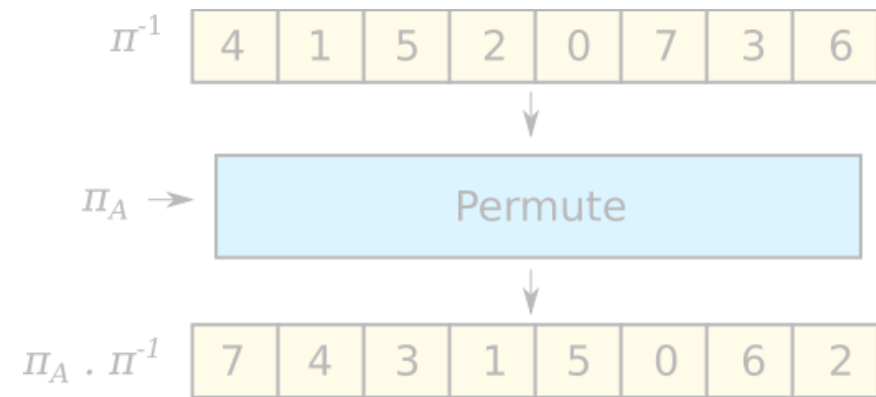
Per-access amortized cost: $\Theta(\sqrt{n} \log n)$




Four-element ORAM



Larger Sizes




4-Block ORAM



$$\begin{aligned}\text{Cost: } & 5B + B + 2B + 3B + \dots \\ & = 11B \text{ every 3 accesses}\end{aligned}$$


Comparison

Linear scan




Cost: $4B = 12B/3$

Our scheme

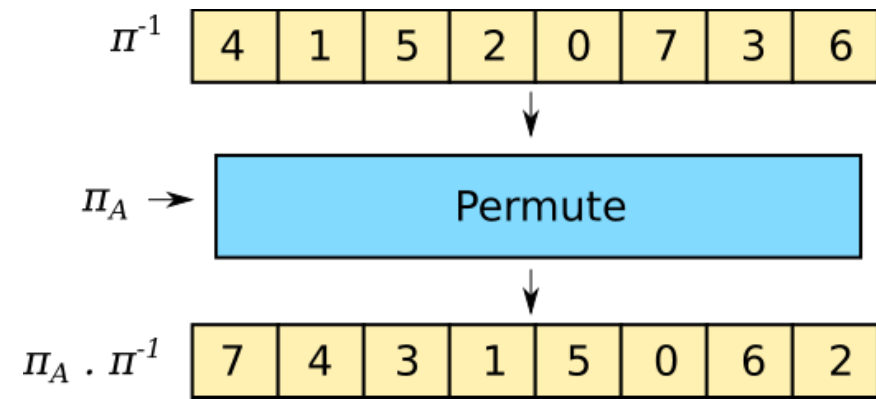


Cost: $11B/3$

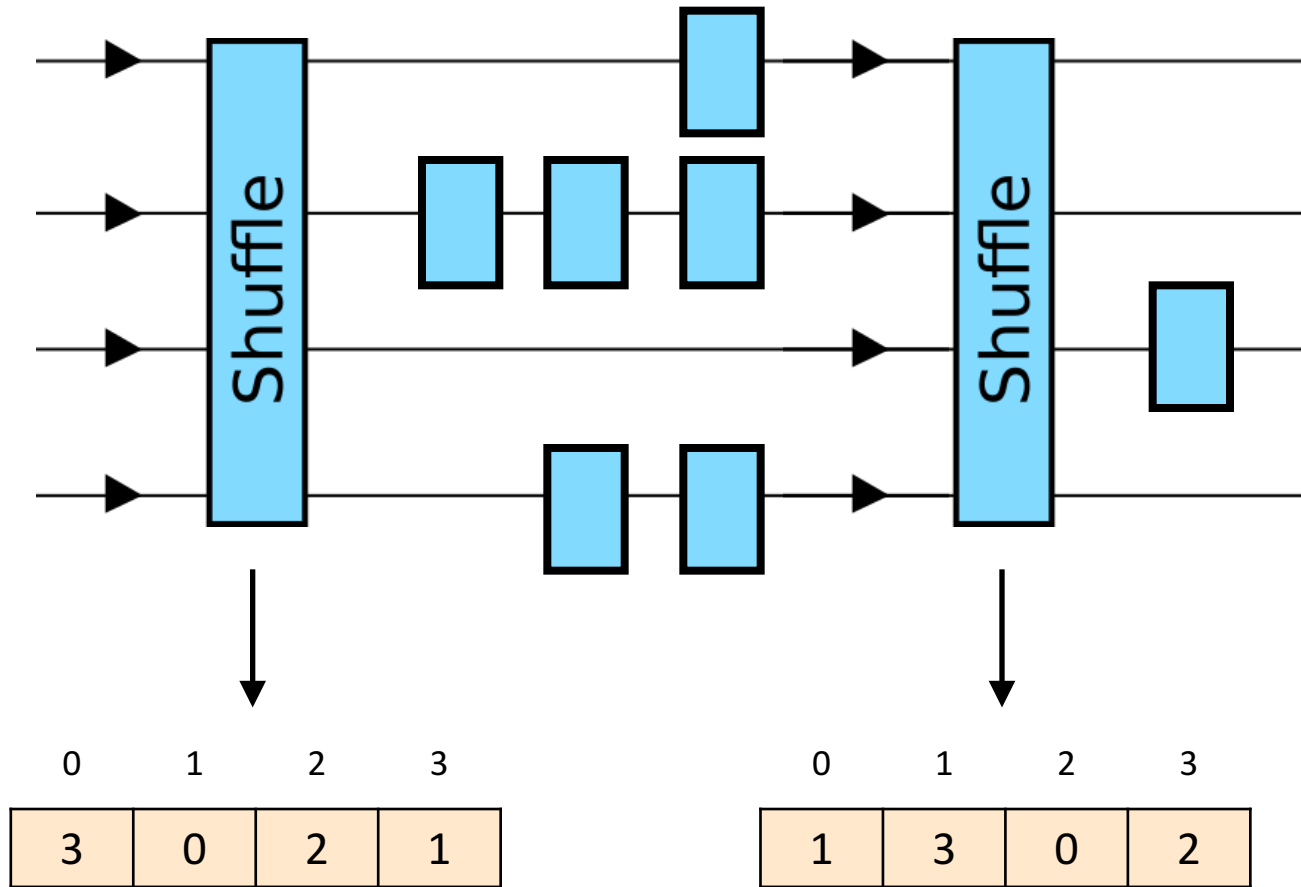
Four-element ORAM



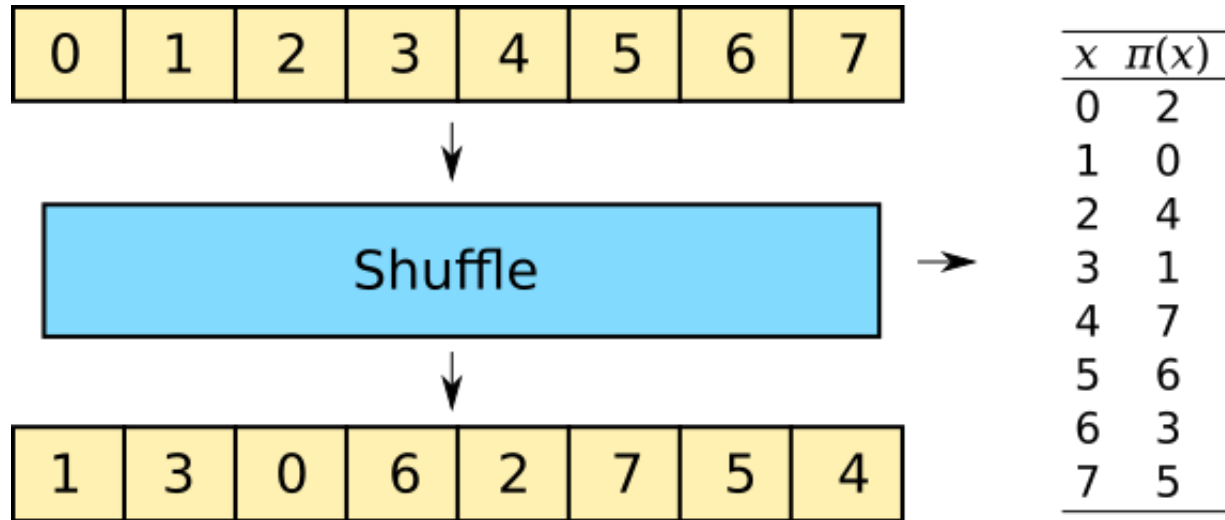
Larger Sizes



Position map



Creating position map



Creating position map

0	1	2	3	4	5	6	7
---	---	---	---	---	---	---	---



Shuffle



1	3	0	6	2	7	5	4
---	---	---	---	---	---	---	---



Shuffle



4	1	5	2	0	7	3	6
---	---	---	---	---	---	---	---

x	$\pi(x)$
0	4
1	1
2	3
3	6
4	0
5	2
6	7
7	5

Inverse permutation



© www.21CARTOONDRAWING5.COM

π_A

4	1	5	2	0	7	3	6
---	---	---	---	---	---	---	---

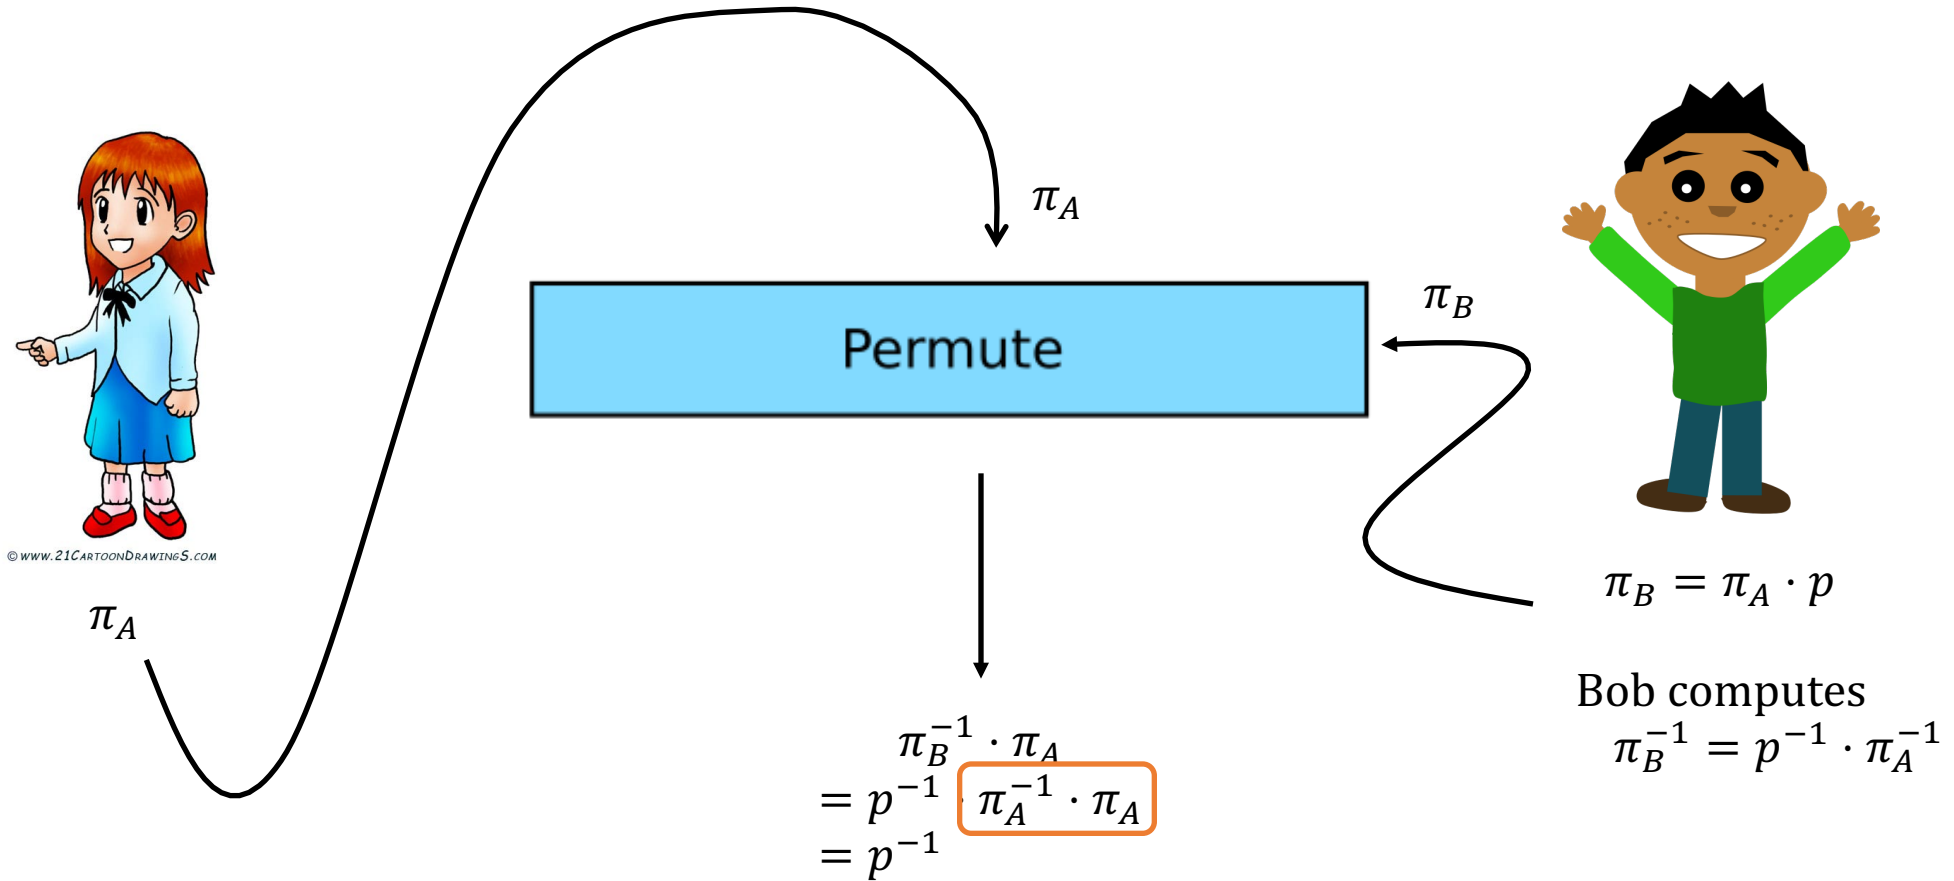
↓ p

Permute



$\pi_B = \pi_A \cdot p$

Inverse permutation




Rinse and repeat




1. Shuffle elements
2. Recreate position map
3. Service $T = \sqrt{n \log n}$ accesses

Access time



Initialization cost



Benchmarks

Task	Parameters	Linear scan	Circuit ORAM	Square-root ORAM
Binary search	2^{10} searches 2^{15} elements	1020 s	5041 s	825 s
Breadth-first search	2^{10} vertices 2^{13} edges	4570 s	3750 s	680 s
Stable matching	2^9 pairs	-	189000 s	119000 s
script hashing	$N = 2^{14}$	≈ 7 days	2850 s	1920 s

Conclusion

We revisited a well-known scheme and used it to

- Lower initialization cost
- Improve breakeven point

Shows that asymptotic costs are not the final word, concrete costs require more consideration.

Download

oblivc.org/sqoram

Contact for help:

Samee Zahur <samee@virginia.edu>