From Mercury Delay Lines to Magnetic Core Memories: **Progress in Oblivious Memories**

David Evans University of Virginia <u>www.cs.virginia.edu/evans</u> <u>oblivc.org</u>

Theory and Practice of Secure Multiparty Computation 2016 Aarhus University 1 June 2016



Setting

Semi-honest model

Two-party computation

Mostly standard assumptions

(although implementation uses Free-XOR)

Oblivious Data Structures



Samee Zahur and David Evans. *Circuit Structures* for Improving Efficiency of Security & Privacy Tools. IEEE Security and Privacy (Oakland) 2013.

Crazy Things in Typical Code





Easy (and Common) Case







Locality: Stacks and Queues if (x != 0)a[i] += 1 if (a[i] > 10)Data-oblivious code i += 1 No branching allowed a[i] = 5

t := a.top() + 1a.cond_update(x != 0, t) $a.cond_push(x != 0 \&\& t > 10, *)$ a.cond_update(x != 0, 5)





Naïve Conditional Push



9

Naïve Conditional Push



10

More Efficient Stack



Block size = 2^{level} Each level has 5 blocks, at least 2 full and 2 empty





Efficient queue operations







Spatial Locality

Not just for stacks and queues

Access cost $\Theta(\log n)$





Temporal Batching

 $\Theta(n \log^2 n)$

Example Application: **DBScan**

Density-based clustering: depth-first search to find dense clusters



Alice's Data Bob's Data **Joint Clusters**

Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu. KDD 1996







15

```
n \leftarrow |P|
c \leftarrow 0
s \leftarrow \text{emptyStack}
cluster \leftarrow [0, 0, \ldots]
for i \leftarrow [1, n] do
    if cluster[i] \neq 0 then
         continue
    V \leftarrow \text{getNeighbors}(i, P, minpts, radius)
    if count(V) < minpts then
         continue
                                          ▷ Start a new cluster
    c \leftarrow c + 1
    for j \leftarrow [1, n] do
         if V[j] = \mathbf{true} \wedge cluster[j] \neq 0 then
              cluster[j] \leftarrow c
              s.push(j)
                                          Conditional Push!
    while s \neq \emptyset do
         k \leftarrow s.pop()
         V \leftarrow \text{getNeighbors}(k, P, minpts, radius)
         if count(V) < minpts then
             continue
         for j \leftarrow [1, n] do
                                                              ⊳ (C)
             if V[j] = true \wedge cluster[j] \neq 0 then
                  cluster[j] \leftarrow c
                                                     Array update!
                  s.push(j)
```

Private Input: *P*-array of points (combines private points from both parties) **Public inputs:** *minpts, radius* **Output:** cluster number for each point



Optimized Structures

55 minutes

480

17

Data-Oblivious Memory

Specialized memory access

Circuit structures, protocol agnostic

Stacks, queues, batched map operations

But first...Obliv-C

General random access Oblivious RAM

Tools for Building Secure Computations

Library-based frameworks: Circuit-level programs

Full control Low-level programming Little type safety



Little control High-level programming Strong type safety

Tools for Building Secure Computations

Library-based frameworks: Circuit-level programs

Full control Low-level programming Little type safety

High-level programming Low-level customizability Helpful, escapable type checking **High-level** Languages

Little control High-level programming Strong type safety

Obliv-C

```
#include<obliv.ob>
#include"million.h"
void millionaire(void* args)
  protocolIO *io=args;
  obliv int v1, v2;
  bool eq,lt;
  v1 = feedOblivInt(io->mywealth,1);
  v2 = feedOblivInt(io->mywealth,2);
  obliv bool eqx = (v1==v2);
  obliv bool ltx = (v1<v2);
 revealOblivBool(&eq,eqx,0);
  revealOblivBool(&lt,ltx,0);
  io->cmp = (!eq?lt?-1:1:0);
}
```

Obliv-C

```
#include<obliv.oh>
#include"million.h"
```

```
void millionaire(void* arc
```

```
protocolIO *io=args;
obliv int v1, v2;
bool eq,lt;
```

```
v1 = feedOblivInt(io->my
v2 = feedOblivInt(io->my
obliv bool eqx = (v1==v2
obliv bool ltx = (v1 < v2)
revealOblivBool(&eq,eqx,
revealOblivBool(&lt,ltx,
io - cmp = (!eq?lt? - 1:1:0)
```

#include <million.h>

```
int main (int argc, char *argv[]) {
   ProtocolDesc pd;
   ProtocolIO io;
   int p = (argv[1] == '1' ? 1 : 2);
   sscanf(argv[2], "%d", &io.myinput);
   // ... set up TCP connections
   setCurrentParty(&pd, p);
   execYaoProtocol(&pd, millionaire, &io);
  printf("Result: %d\n", io->cmp);
   // ... cleanup }
```

Oblivious Conditionals

```
obliv if (cmpresult < 0) {
    iimin = iimid + 1;
} else {
    iimax = iimid;
}</pre>
```

from <a>obinary_search

```
obliv int obinary_search_oram(OcCopy * cpy, void* result, oram * haystack, void* needle, block_cmp_function fn) {
       int upper bound = log2(oram size(haystack)) + 1;
```

```
obliv int index = -1;
obliv int iimin = 0;
obliv int iimax = oram_size(haystack) - 1;
obliv int iimid;
obliv int cmpresult;
void * temp element = calloc(1, cpy->eltsize);
for (int ii = 0; ii < upper_bound; ii++) {</pre>
        iimid = (iimin + iimax) / 2;
        oram read(temp element, haystack, iimid);
        cmpresult = fn(cpy, temp_element, needle);
        obliv if (cmpresult == 0) {
                ocCopy(cpy, result, temp element);
                index = iimid;
        } else {
                obliv if (cmpresult < 0) {</pre>
                         iimin = iimid + 1;
                } else {
                         iimax = iimid;
                }
        }
free(temp element);
return index;
```

}

Actual code...with all the ugly parts

Escaping

{ ~obliv(var)

Code inside ~obliv always executes regardless of oblivious condition

var is Boolean: oblivious condition

Programmer has control! But, not security risk: all private data is still encrypted

```
#include <math.h>
#include "oqueue.oh"
                                      level 0
struct oqueue {
        oqueue * child;
                                      t=2
        OcCopy * cpy;
        size_t layer_index;
                                      level 0
        size_t copymult;
        bool dynamic_sizing;
                                       h=5
        obliv uint8_t head;
        obliv uint8_t tail;
        obliv uint32_t current_elements;
        uint8_t push_time;
        uint8_t pop_time;
        void * data;
};
```

Implementing Oblivious Queue



```
obliv bool oqueue_push(oqueue * layer, void * input) obliv {
        obliv bool success = false;
        ~obliv(en)
                if (layer->child != NULL) {
                        if (layer->push time == 1) {
                                 obliv if (layer->tail >= 5) {
                                         obliv if (oqueue_push(layer->child, element(layer->cpy,layer->data,3*layer->copymult))) {
                                                 ocCopyN(layer->cpy,element(layer->cpy,layer->data,3*layer->copymult), element(layer-
                                                 layer->tail -= 2;
                                         }
                                 layer->push time = 0;
                        } else {
                                 layer->push time ++;
                        }
                }
        obliv if (layer->tail < 6) {</pre>
                for (uint8_t ii = 0; ii < 6; ii++) {</pre>
                        obliv if (ii == layer->tail) {
                                 ocCopyN(layer->cpy,element(layer->cpy,layer->data,ii * layer->copymult), input, layer->copymult);
                                 success = true;
                                 layer->current_elements ++;
                        }
                layer->tail++;
        return success;
```



http://oblivc.org/

Obliv-C Tutorial

Obliv-C Tutorial Home

Obliv-C Documentation

Introduction to Obliv-C

NOTE: This webpage is currently being written and is subject to change frequently. Obliv-C is a lightweight GCC wrapper that makes it simple for multiple parties to compute a joint function on their private inputs. A full description of the language can be found here. The language comes with numerous built-in protocols which allow a programmer to quickly implement a secure multiparty computation. One particularly useful application of Obliv-C is in allowing for aggregate data analysis from private datasets. Check out the documentation on various built-in functionality.



Historical Excursion

A Permutation Network

ABRAHAM WAKSMAN

Stanford Research Institute, Menlo Park, California

ABSTRACT. In this paper the construction of a switching network capable of n!-permutation of its n input terminals to its n output terminals is described. The building blocks for this network are binary cells capable of permuting their two input terminals to their two output terminals.

The number of cells used by the network is $\langle n \cdot \log_2 n - n + 1 \rangle = \sum_{k=1}^n \langle \log_2 k \rangle$. It could be argued that for such a network this number of cells is a lower bound, by noting that binary decision trees in the network can resolve individual terminal assignments only and not the titioning of the permutation set itself which requires only $\langle \log_2 n! \rangle = \langle \sum_{k=1}^n \log_2 k \rangle$ binary `ans.

January 1968

Journal of the ACM,

1967 ACM Turing Lecture

Computers Then and Now

MAURICE V. WILKES

Cambridge University, Cambridge, England

(In same Jan 1968 JACM as Waksman Network!)



	READ	WRITE	RECIRCULATE
NEAD SELECT	1	0.041	D CH 1
WRITE SELECT	0.081	1	0
DATA INPUT	0 OR 1	0 OR 1	0 OR 1
R/W ENABLE	1	1	0 OR1



Fro. 2.1.1. Delay line memory block diagram.

Delay Lines



Mercury Delay Lines



Why Mercury?

Speed of Sound 343 m/s Air Mercury 1450 m/s (40° C) Water 1500 m/s (25° C)



Turing's contribution to this discussion was to advocate the use of gin, which he said contained alcohol and water in just the right proportions to give a zero temperature coefficient of propagation velocity at room temperature.



where we are now going unrough.

In ultrasonic memories, it was customary to store up to 32 words end to end in the same delay line. The pulse rate was fairly high, but people were much worried about the time spent in waiting for the right word to come around. Most delay line computers were, therefore, designed so that, with the exercise of cunning, the programmer could place his instructions and numbers in the memory in such a way that the waiting time was minimized. Turing himself was a pioneer in this type of logical design. Similar methods were later applied to computers which used a magnetic drum as their memory and, altogether, the subject of optimum coding, as it was called, was a flourishing one. I felt that this kind of human ingenuity was misplaced as a long-term investment, since sooner or later we would have truly random-access memories. We therefore did not have anything to do with optimum coding in Cambridge.

Magnetic

Core



MIT Project Whirlwind, 1951 2K 16-bit words with "no waiting"!
Oblivious RAM

, and gooner, one subject of optimum coding, as it was called, was a flourishing one. I felt that this kind of human ingenuity was misplaced as a long-term investment, since sooner or later we would have truly random-access memories. We therefore did not have anything to do with optimum coding in Cambridge.

Linear Scan Doesn't Scale

Writing a single 32-bit integer: 32 logic gates Raw Yao's performance ≈ 3M gates per second Write speed \approx 100,000 elements per second (not hiding access pattern)

> For hiding access pattern, $N = 2^{17}$ elements requires > 1 second per access



Security property: all initialization and access sequences of the same length are indistinguishable to server.

Linear server-side encrypted state

RAM-SC

[Gordon, Katz, Kolesnikov, Krell, Malkin, Raykova, Vahlis 2012]





Circuit ORAM Access time

State-ofthe-ORAM-Art in 2015

Xiao Wang, Hubert Chan, and Elaine Shi. Circuit ORAM: On Tightness of the Goldreich-Ostrovsky Lower Bound. In ACM CCS 2015.

 $\Theta(\log^3 n)$





(Not including initialization cost) 2^{10} 2^{12} 2^{11}

Classical Square-Root ORAM



Ostrovsky et al.

COMPREHENSIVE SOFTWARE [54] **PROTECTION SYSTEM**

- Inventors: Rafail Ostrovsky, Acton, Mass.; [75] Oded Goldreich, Tel Aviv, Israel
- Massachusetts Institute of Assignee: [73] Technology, Cambridge, Mass.
- Appl. No.: 476,814 [21]
- Feb. 7, 1990 [22] Filed:

Related U.S. Application Data

Continuation-in-part of Ser. No. 395,882, Aug. 18, [63] 1989, abandoned.

[51]	Int. Cl. ⁵	 H0	4L 9	/00
[52]	U.S. Cl.	 380/4;	380/	/46;
			364/	969

Sorting Network," Pr Computing, 1983. N. Pippenger and M **Complexity Measures** puting Machinery, 26(O. Goldreich, S. Go Cryptographic Appli Proc. of CRYPTO-84, M. Luby and C. Racl random Permutations SIAM J. Comp., 17(2) O. Goldreich, S. Gol Construct Random F tion for Computing 792-807. A. V. Aho, J. E. Hop and Analysis of Comp

[11]

[45]



Problems with SQ-ORAM Design

- Requires a PRF for each ORAM access – PRF is a big circuit in MPC
- Initialization requires $N + \sqrt{N}$ PRF evaluations
- Requires $\Theta(Nlog^2N)$ oblivious sort twice:
 - Shuffling memory according to PRF
 - Removing dummy blocks

ution strategy: use random permutation instead of PRF



Shuffling Network [Waksman 1968]



Cost per shuffle: 5*B*







Linear scan



Less expensive than linear scan for 4 blocks (8 with overhead)









)	1	2	3	4	5	6	7	8	9

0	1	2	3	4	5	6	7	8	9
()	-	1		>	5	2		1
 <u>`</u>			•		-	<u>`</u>	-		•

0	1	2	3	4	
(0			2	2

Stash (empty)





Logical index/2



Logical index/4

Stash (empty)



























9 3 0 4 5	1	7	6	2
-----------	---	---	---	---

9	9	1	3	6	8	7	4	5
	1		0		с, С	3	2	2



After First Access

read a[9]



2	0	4	3	
(0		I	

6	2





6	2
---	---



6	2
---	---



6	2
---	---

After Second Access

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











Creating position map



Creating position map



Inverse permutation



Alice picks a random masking permutation

permutation revealed to Bob



 $\pi_B^{-1} = p^{-1} \cdot \pi_A^{-1}$



- $\langle \mathbf{if} \rangle$ Oram₀.Stash_{*i*}. $\langle \mathsf{index} \rangle = \langle i \rangle$: $p \leftarrow \text{GetPos}(\text{Oram}_0.\text{Oram}_1, \langle i \rangle, \langle \text{found} \rangle)$
- **append** Oram₀.Shuffle_p **to** Oram₀.Stash $Oram_0.Shuffle_{p'} \leftarrow Oram_0.Stash_j$
 - $Oram_0 \leftarrow Initialize(Oram_0.Shuffle)$



Pre-Access Cost (not counting initialization)







16-byte blocks 32-byte blocks



16-byte blocks 32-byte blocks
Benchmark	Parameters	Linear Scan	Square-Root ORAM
Binary Search	1 search	1.00	10.41
	2^5 searches	31.87	26.25
	2^{10} searches	1019.77	824.81
Breadth-First Search	$n = 2^2$	0.09	0.34
	$n = 2^5$	4.77	4.08
	$n = 2^{10}$	4569.31	679.63
Gale-Shapley	2^3 pairs	-	0.51
	2 ⁶ pairs	-	145.13
	2 ⁹ pairs	-	119405.
Scrypt	$N = 2^5$	4.11	3.43
	$N = 2^{10}$	1678.16	293.79
	$N = 2^{14}$	about 7 days	1919.92
	Litecoin	210.92	40.29

Wall-clock time in seconds for full protocol between two EC2 C4.2xlarge nodes (1.03 Gbps)

Circuit ORAM

3228.69 3282.40 5040.82

4.28 42.66 3750.57

6.57 1328.50 188972.

34.47 1453.85 2846.51 247.29 odes (1.03 Gbps)

Benchmark	Parameters	Linear Scan	Square-Root ORAM
Binary Search	1 search	1.00	10.41
	2^5 searches	31.87	26.25
	2 ¹⁰ searches	1019.77	824.81
Breadth-First Search	$n = 2^2$	0.09	0.34
	$n = 2^5$	4.77	4.08
	$n = 2^{10}$	4569.31	679.63
Gale-Shapley	2^3 pairs	-	0.51
	2 ⁶ pairs	-	145.13
	2 ⁹ pairs	-	¹¹ ~32 minute
Scrypt	$N = 2^{5}$	4.11	55,000 <i>x</i> sta
	$N = 2^{10}$	1678.16	793.19
	$N = 2^{14}$	about 7 days	1919.92
	Litecoin	210.92	40.29

Wall-clock time in seconds for full protocol between two EC2 C4.2xlarge nodes (1.03 Gbps)

Circuit ORAM

3228.69 3282.40 5040.82

4.28 42.66 3750.57

6.57 1328.50

es andard execution 1453.85 2846.51 247.29 nodes (1.03 Gbps)

Benchmark	Parameters	Linear Scan	Square-Root ORAM
Binary Search	1 search	1.00	10.41
	2^5 searches	31.87	26.25
	2^{10} searches	1019.77	824.81
Breadth-First Search	$n = 2^2$	0.09	0.34
	$n = 2^5$	<u> 1</u> 77	4 08
	$n = 2^1 \sim 33$ hours ("wikipedia" version)		
Gale-Shapley	2 ³ pai Imp	proved to ~1 h	our with custom st
	2^6 pairs	-	145.13
	2 ⁹ pairs	-	119405.
Scrypt	$N = 2^5$	4.11	3.43
	$N = 2^{10}$	1678.16	293.79
	$N = 2^{14}$	about 7 days	1919.92
	Litecoin	210.92	40.29

Wall-clock time in seconds for full protocol between two EC2 C4.2xlarge nodes (1.0 Gbps)

Circuit ORAM

3228.69 3282.40 5040.82

4.28

- 47 66
- 57

ructures 57

1328.50

188972.

34.47 1453.85 2846.51 247.29 nodes (1.0 Gbps)

Open Problems

- **Scalability:** poly-logarithmic hierarchical ORAM design
- Automatic optimization: using custom data structures when memory access predictable
- Stronger security models: active security

All results are semi-honest model

Establishing Meaningful Trust



64 KB memory $1 \,\mu s$ access (~2000x improvement)

Collaborators







Samee Zahur Jack Doerner David Evans

Xiao Wang Jonathan Katz Mariana Raykova

Code and Paper: oblivc.org/sqoram

Adrià Gascón



David Evans evans@virginia.edu





www.cs.virginia.edu/evans OblivC.org mightBeEvil.org