#### **PIR-PSI: SCALING PRIVATE CONTACT DISCOVERY**

**PETS 2018** 



#### Daniel Demmler

Peter Rindal Mike Rosulek Ni Trieu



# Motivation – Application: Contact Discovery

- Contact discovery tells social network users which of their contacts are in the social network
- An **insecure** naïve hashing-based protocol is used in practice
- Vulnerable to
  - Brute-force attacks (for small input domain, e.g. phone numbers)
  - Comparison with hashes from later sessions





# Motivation – Application: Private Contact Discovery

- Contact Discovery should be efficient and scalable, and protect the privacy of user inputs.
- It runs once when a user initially joins a social network
- ... and periodically to find contacts that join the social network later on.



# **Private Set Intersection (PSI)**



#### **Private Set Intersection (PSI)**



 $X \cap Y$ 

#### **PSI for Contact Discovery**



#### **Status-Quo vs. PIR-PSI**

- Communication linear in both sets O(N + n)
  - What about  $N \gg n$ ?
  - Insecure solution
    - Send small set to other party
    - Communication = O(min(N, n))



#### • PIR-PSI

- Communication =  $O\left(n\log\left(\frac{N\log n}{n}\right)\right)$
- Client Computation =  $O\left(n\log\left(\frac{N\log n}{n}\right)\right)$  AES operations
- Server Computation = O(N log n) AES operations

#### **Plaintext Database Query**



#### **Private Information Retrieval (PIR)**



#### 2-Server PIR [CGKS95]



#### 2-Server PIR [CGKS95]



#### Example: 2-Server Linear Summation PIR [CGKS95]



# **PIR from Distributed Point Functions (DPFs)**

- Point Functions:  $PF = \{f_{i,v}: f_{i,v}(i) = v, f_{i,v}(x) = 0 \forall x \neq i\}.$
- **Distributed** PFs allow 2 parties the secretshared PF evaluation, without revealing *i*, *v*.
- DPFs are described by short keys k<sub>1</sub>, k<sub>2</sub> of length O(log N), where N is the domain of i.
- By using v = 1, i.e., a DPF returning 1 only at index i, we can express the plain text query q and thus build 2-server PIR with O(log N) communication complexity.
- Instantiated efficiently with AES.



#### **Designated-Output PIR**



#### **PIR Private Equality Test**













$$\leftarrow \text{ Collision: } h(y_1) = h(y_N)$$



- To avoid collisions: use multiple hash functions in this example: *h*, *h*'.
- In our implementation we used 3 hash functions and a cuckoo expansion factor of  $e \approx 1.4$  for a cuckoo failure probability of  $2^{-20}$  during one-time initialization.

• Every element can be located in two possible bins.



• The client computes all hash functions for every element.

• Every element can be located in two possible bins.



- To check if the server holds  $x_1$ , the client runs a PIR-PEQ with the 2<sup>nd</sup> and 4<sup>th</sup> bin.
- In the full protocol: instead of single PIR-PEQ, we run all of them together in a PSI protocol.

#### **PIR-PSI Overview**

- 1. Cuckoo Hashing
  - Both servers compute the same cuckoo hash table for their *N* elements.
- 2. DPF-PIR Query
  - The client delegated extraction of *n* elements from the cuckoo table.

#### 3. Oblivious Shuffle

• One server receives the other server's masked output and obliviously shuffles the PIR results to hide which Cuckoo hash function was used.

#### 4. Small PSI

• A standard PSI protocol is used to determine intersection.



# Optimizations

#### Binning

- Instead of running full domain DPFs, we partition the server cuckoo table into bins and a smaller DPFs per bin.
- Parallelization!
- Batching
  - Instead of running DPF queries separately, run all queries in each bin in parallel.
  - Only a single pass over the cuckoo table for multiple queries.
- Larger PIR Blocks
  - PIR queries can return multiple cuckoo table entries.
  - less communication, more computation in PSI.







#### **PIR-PSI with 3 PIR Servers**



$$(K_{2} \cdot DB_{2}) \bigoplus (K_{2} \cdot DB_{3}) \bigoplus (K_{1} \cdot DB) \bigoplus m$$

$$=$$

$$K_{2} \cdot (DB_{2} \bigoplus DB_{3}) \bigoplus (K_{1} \cdot DB) \bigoplus m$$

$$=$$

$$K_{2} \cdot DB \bigoplus K_{1} \cdot DB \bigoplus m$$

$$=$$

$$DB[i] \bigoplus m$$

#### **PIR-PSI Performance**

- Communication and running time for  $n = 1\ 024$  client elements and server set sizes  $N \in \{2^{20}, 2^{24}, 2^{26}, 2^{28}\}.$
- Benchmarked in Gigabit LAN, on 1 machine with 36 x 2.3 GHz.
   Implementation set to use 4 threads.
- Client computation is  $\approx 10\%$  of total.
- Parameters for communication / computation trade-off

··×· 2<sup>2</sup>0 ···• 2<sup>2</sup>24 ··•• 2<sup>2</sup>26 ··•• 2<sup>2</sup>8



# Conclusion

- Combination of DPF-based PIR with state-of the art PSI to achieve scalable contact discovery.
- Efficient open-source C++ implementation on Github: github.com/osu-crypto/libPSI



- Many more details in the paper!
  - Security Analysis
  - Cuckoo Hashing Parameters
  - Detailed performance analysis and comparison with related work
  - Extensions

# Thank you!



#### Daniel Demmler

Peter Rindal Mike Rosulek Ni Trieu





• Some icons are made by Freepik from flaticon.com

• Extra / Backup slides coming up next...

#### A Sampling of PSI Over the Decades



#### A Sampling of PSI Over the Decades











## A Sampling of PSI Over the Decades

