

Private Set Intersection

CS 598 DH

Today's objectives



Define the private set intersection (PSI) problem

Construct PSI from MPC and discuss circuit PSI

Define Oblivious PRF

Construct PSI for large sets


PSI



{13, 17, 25, 45, 52, 101}


{1, 4, 17, 19, 21, 45, 100}

PSI




{13, **17**, 25, **45**, 52, 101}

{1, 4, **17**, 19, 21, **45**, 100}




{17, 45}

PSI



{13, **17**, 25, **45**, 52, 101}

{1, 4, **17**, 19, 21, **45**, 100}




{17, 45}


Special case of MPC

“Just use MPC”

Because it is a special case, we can hope for much more efficiency



PSI



{13, **17**, 25, **45**, 52, 101}

{1, 4, **17**, 19, 21, **45**, 100}

Efficient Circuit-based PSI via Cuckoo Hashing

Benny Pinkas¹, Thomas Schneider², Christian Weinert², and Udi Wieder³

¹ Bar-Ilan University
benny@pinkas.net

² TU Darmstadt

{thomas.schneider, christian.weinert}@crisp-da.de


³ VMware Research
udi.wieder@gmail.com

Abstract. While there has been a lot of progress in designing efficient custom protocols for computing Private Set Intersection (PSI), there has been less research on using generic Multi-Party Computation (MPC) protocols for this task. However, there are many variants of the set intersection functionality that are not addressed by the existing custom PSI solutions and are easy to compute with generic MPC protocols (e.g., comparing the cardinality of the intersection with a threshold or measuring ad conversion rates).

Generic PSI protocols work over circuits that compute the intersection. For sets of size n , the best known circuit constructions conduct $O(n \log n)$ or $O(n \log n / \log \log n)$ comparisons (Huang et al., NDSS'12 and Pinkas et al., USENIX Security'15). In this work, we propose new circuit-based protocols for computing *variants of the intersection* with an almost linear number of comparisons. Our constructions are based on new variants of Cuckoo hashing in two dimensions.

We present an asymptotically efficient protocol as well as a protocol with better concrete efficiency. For the latter protocol, we determine the required sizes of tables and circuits experimentally, and show that the run-time is concretely better than that of existing constructions.

The protocol can be extended to a larger number of parties. The proof technique presented in the full version for analyzing Cuckoo hashing in



{17, 45}

Special case of MPC


“Just use MPC”

Because it is a special case, we can hope for much more efficiency

PSI



$\{x_0, x_1, x_2, x_3\}$




$\{y_0, y_1, y_2, y_3\}$

PSI



$\{x_0, x_1, x_2, x_3\}$



$\{y_0, y_1, y_2, y_3\}$




$H(y_0), H(y_1), H(y_2), H(y_3)$

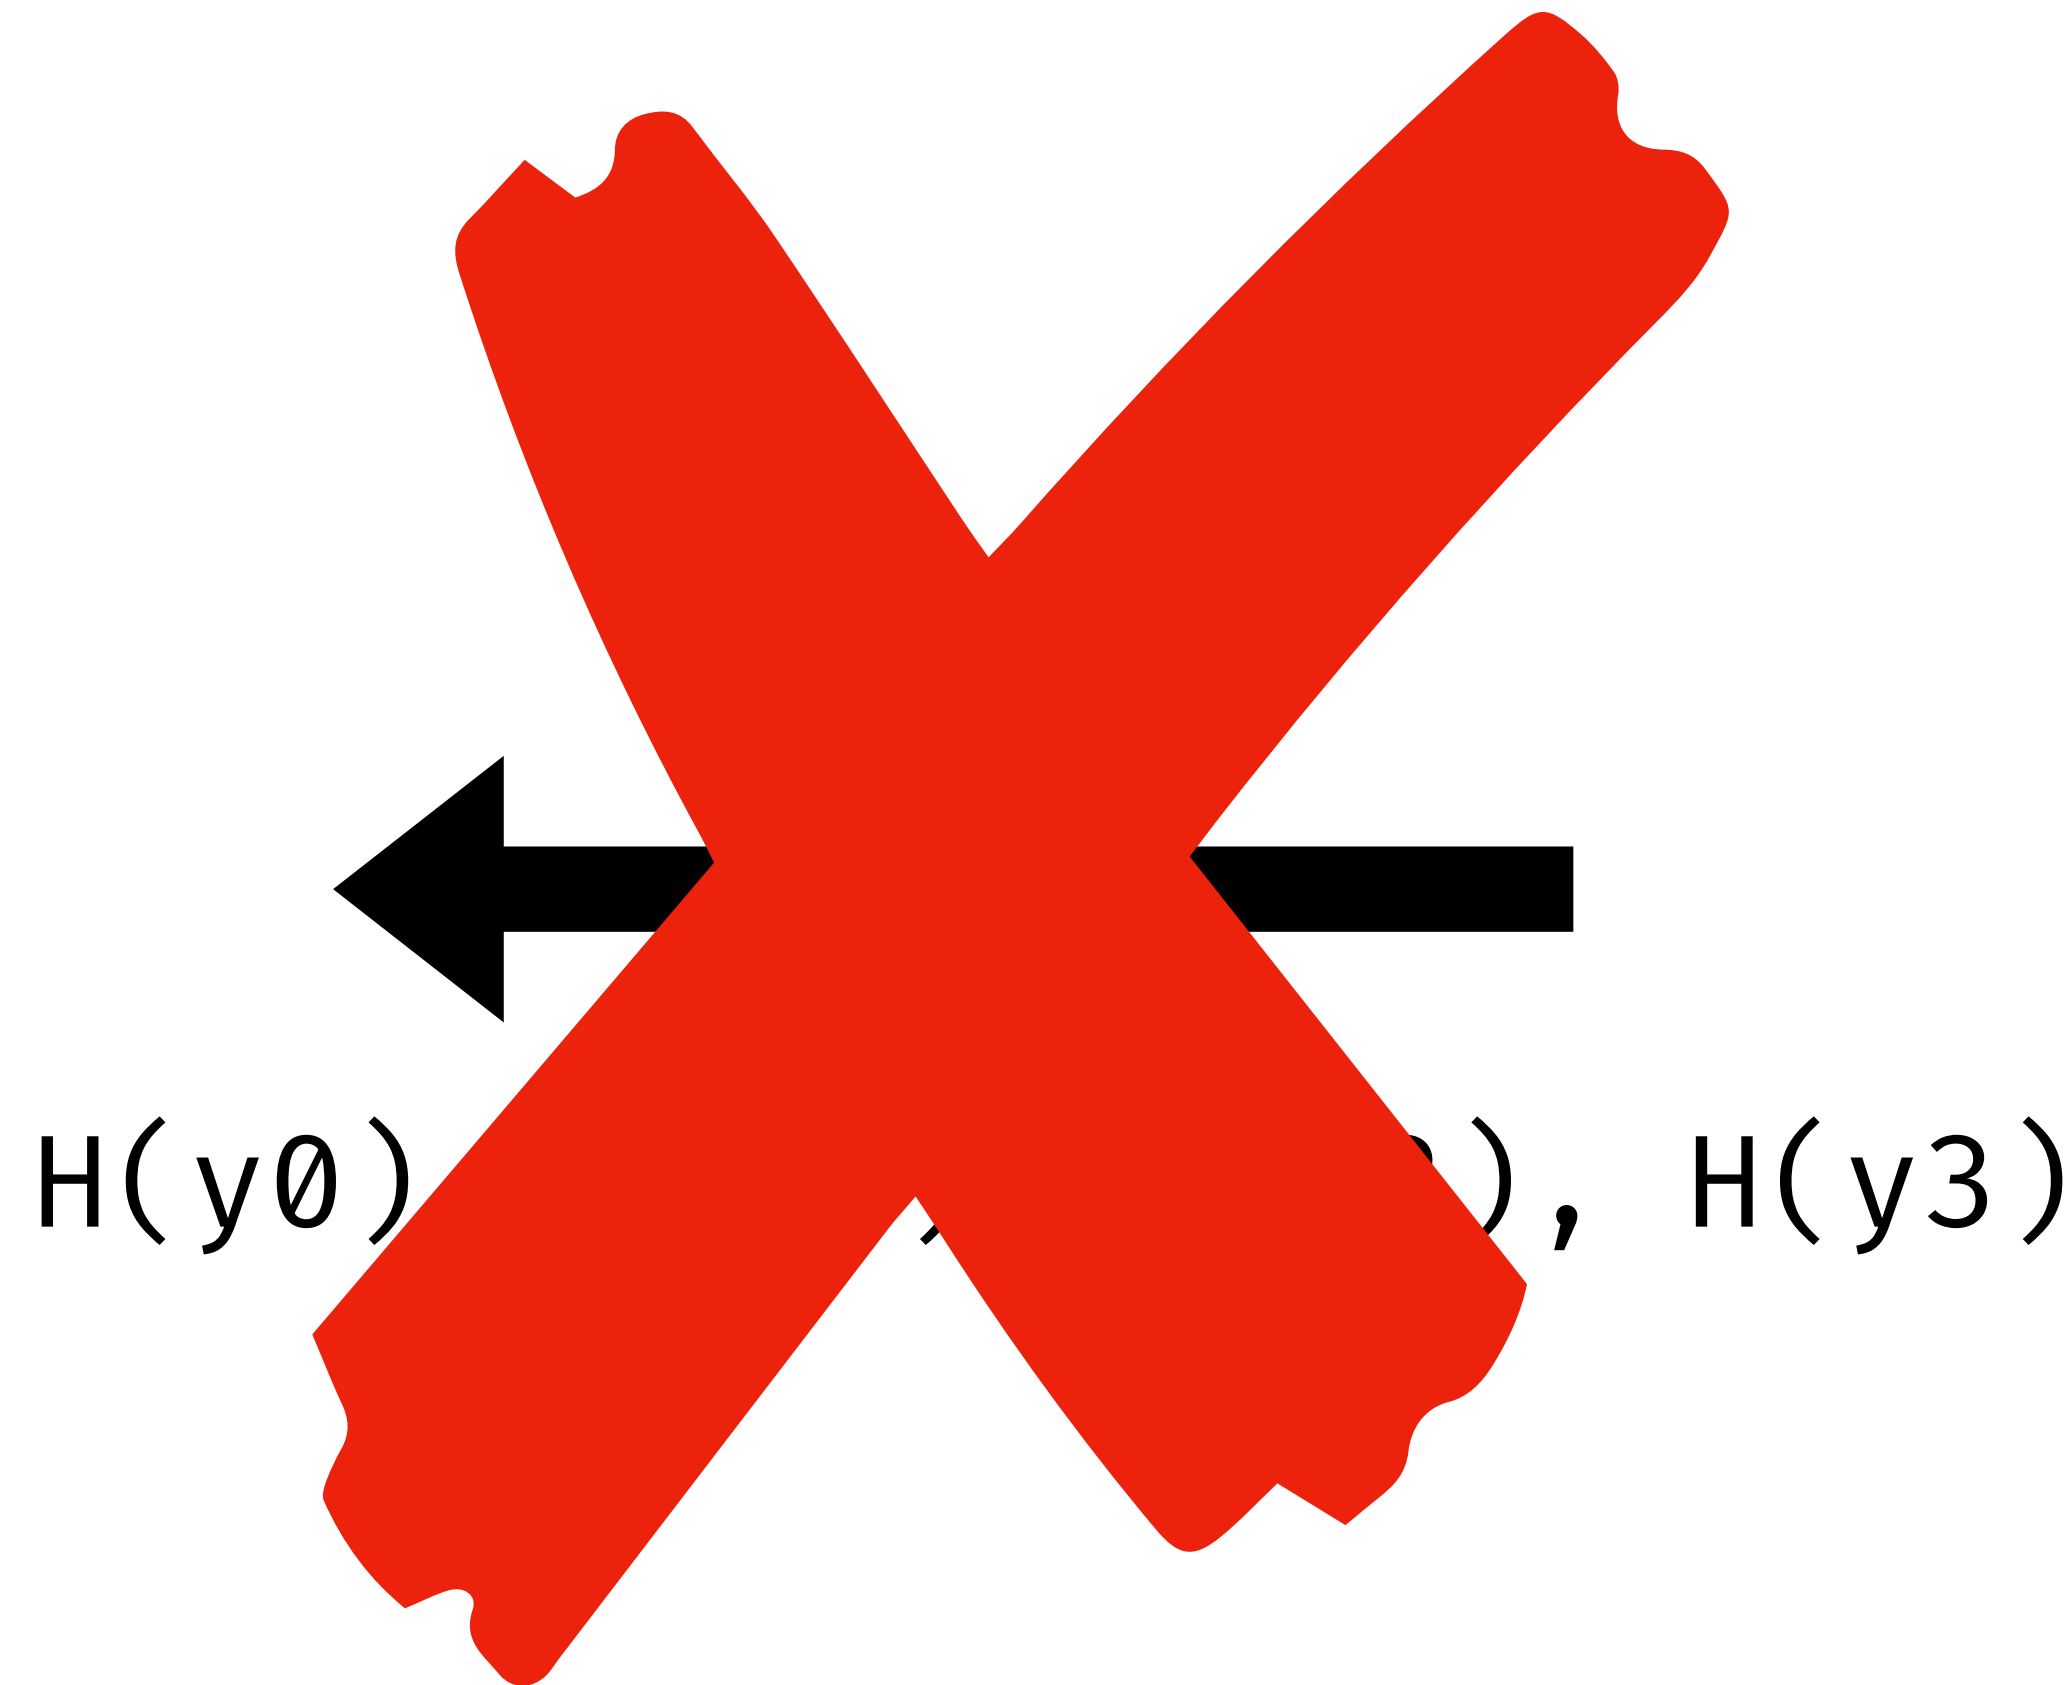
PSI



$\{x_0, x_1, x_2, x_3\}$



$\{y_0, y_1, y_2, y_3\}$




We cannot simulate semi-honest Alice

Oblivious PRF



Oblivious PRF




$\{x_0, x_1, x_2, x_3\}$




$\{y_0, y_1, y_2, y_3\}$



$F(k, x_0), F(k, x_1), \dots$



Batched Oblivious PRF



x_0

0

x_1

1

x_2

2

x_3

3



x_4

4

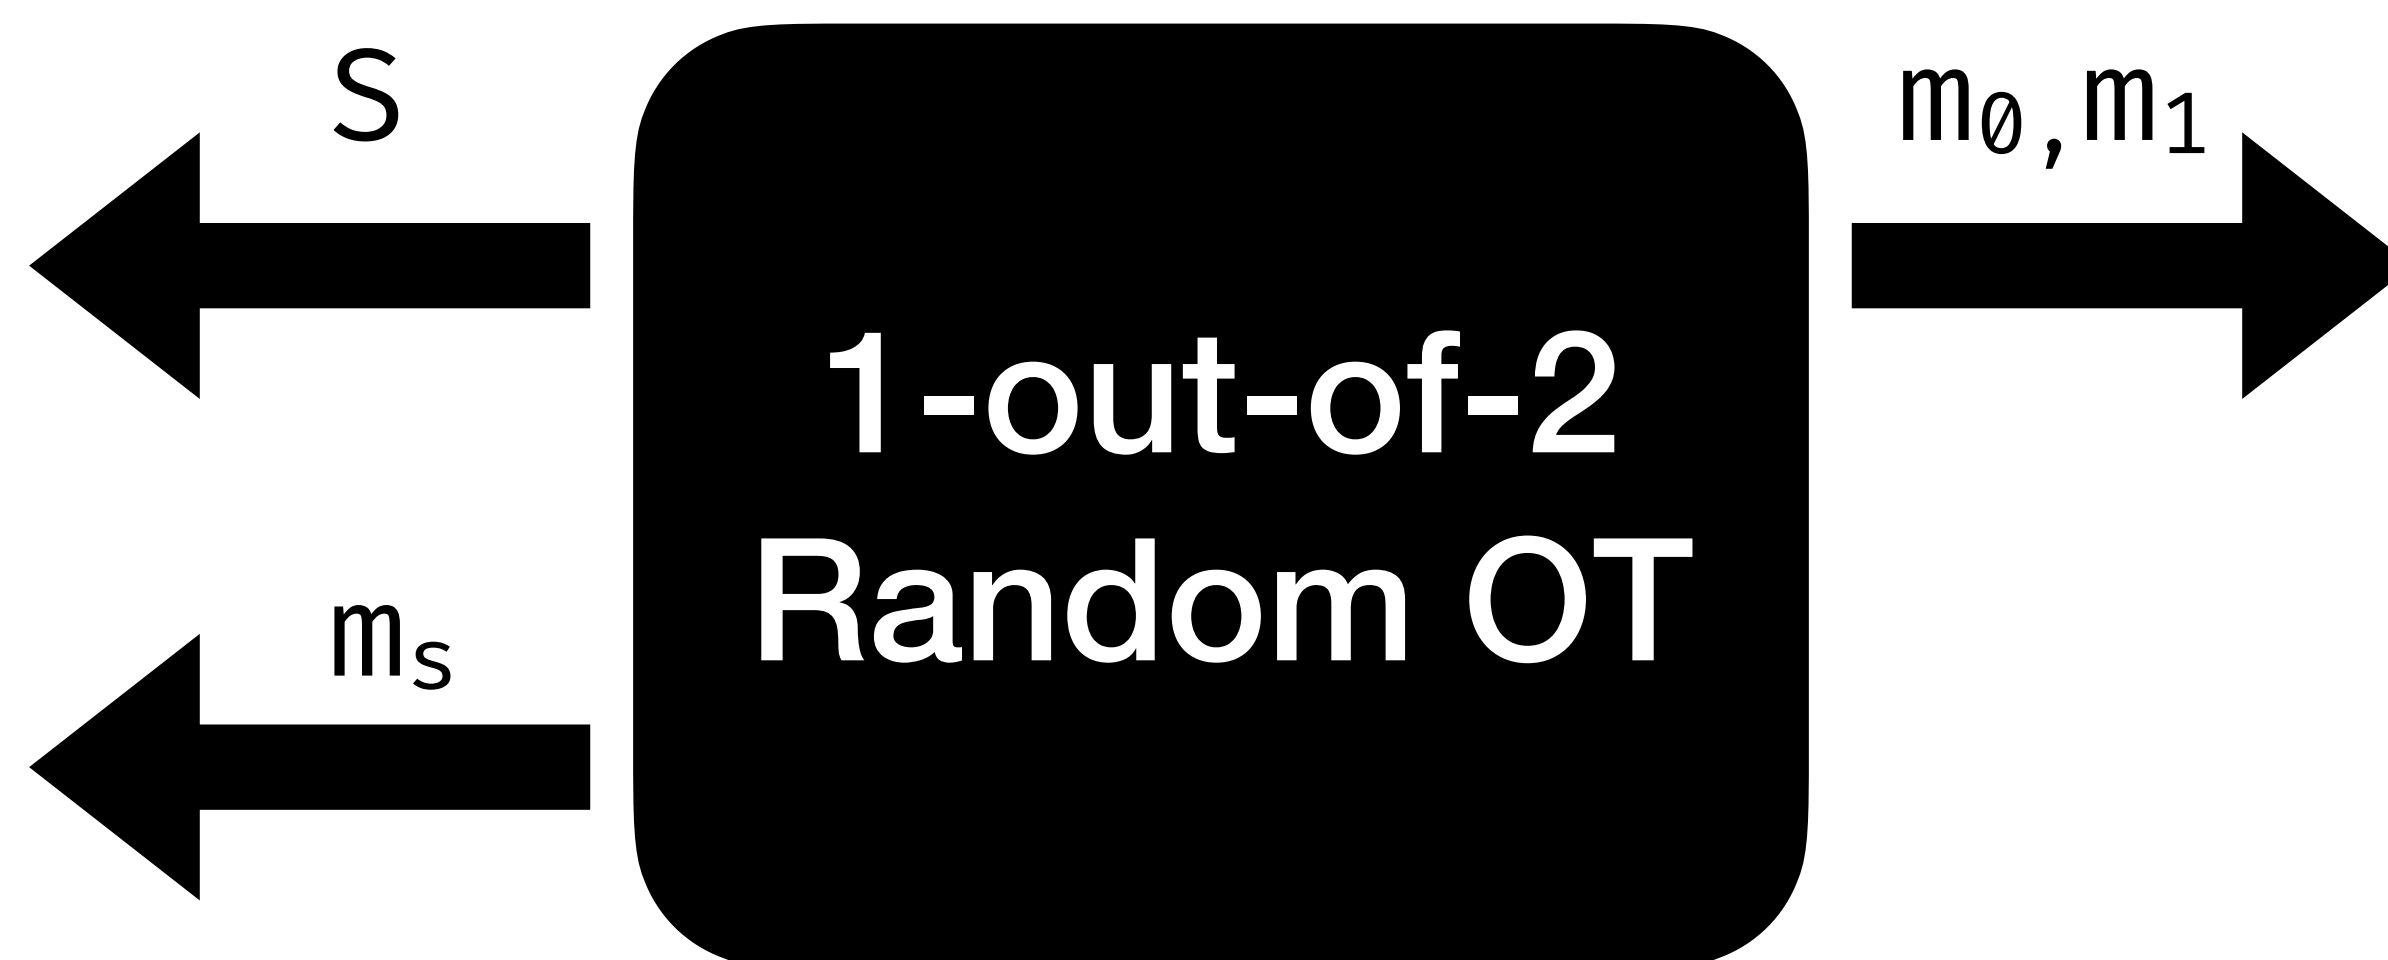

x_5



5


Batched Oblivious PRF




x_0	$F(k_0, x_0)$	0	k_0
x_1	$F(k_1, x_1)$	1	k_1
x_2	$F(k_2, x_2)$	2	k_2
x_3	$F(k_3, x_3)$	3	k_3
x_4	$F(k_4, x_4)$	4	k_4
x_5	$F(k_5, x_5)$	5	k_5








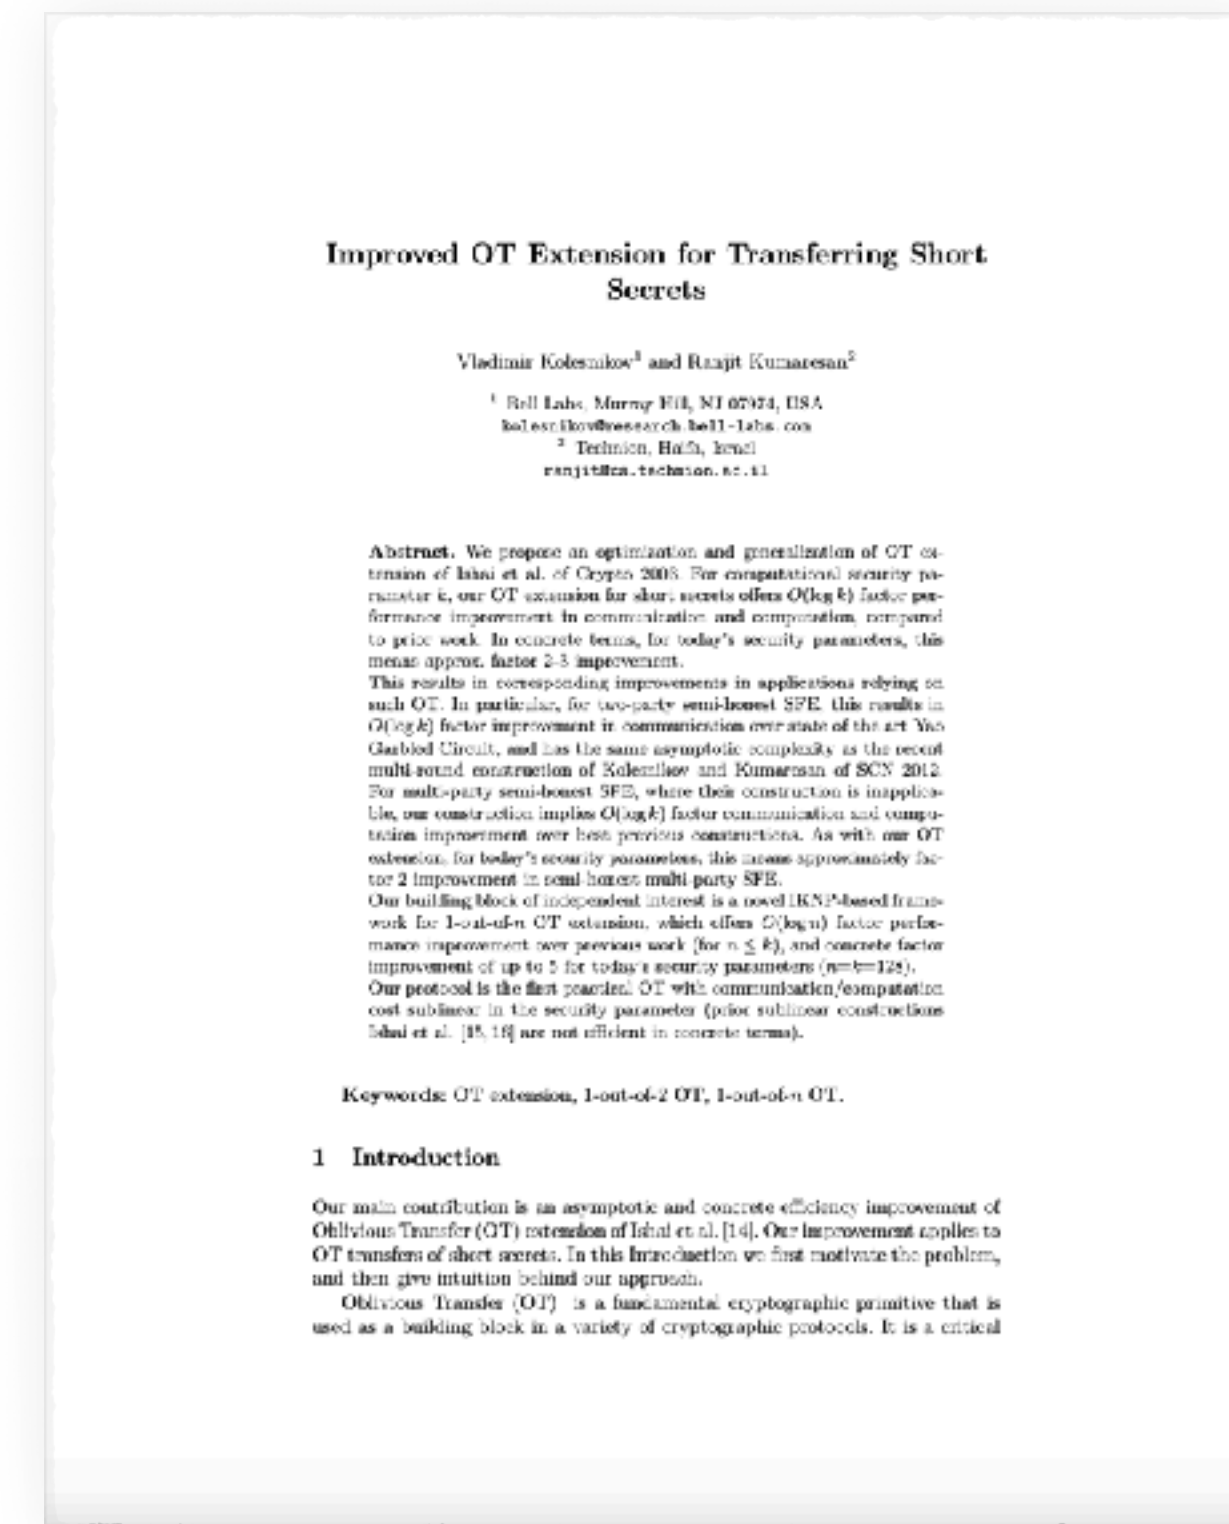


Use $\log n$ OTs




Use $\log n$ OTs

... or do something smarter





Batched Oblivious PRF



Essentially batched 1-out-of-N OT

x_0	$F(k_0, x_0)$	0	k_0
x_1	$F(k_1, x_1)$	1	k_1
x_2	$F(k_2, x_2)$	2	k_2
x_3	$F(k_3, x_3)$	3	k_3
x_4	$F(k_4, x_4)$	4	k_4
x_5	$F(k_5, x_5)$	5	k_5

Efficient Batched Oblivious PRF with Applications to Private Set Intersection

Vladimir Kolesnikov* Ranjit Kumaresan† Mike Rosulek‡ Ni Trieu‡

August 20, 2016

Abstract

We describe a lightweight protocol for oblivious evaluation of a pseudorandom function (OPRF) in the presence of semi-honest adversaries. In an OPRF protocol a receiver has an input r ; the sender gets output s and the receiver gets output $F(s, r)$, where F is a pseudorandom function and s is a random seed. Our protocol uses a novel adaptation of 1-out-of-2 OT-extension protocols, and is particularly efficient when used to generate a large batch of OPRF instances. The cost to realize m OPRF instances is roughly the cost to realize $3.5m$ instances of standard 1-out-of-2 OTs (using state-of-the-art OT extension).

We explore in detail our protocol's application to semi-honest secure private set intersection (PSI). The fastest state-of-the-art PSI protocol (Pinkas et al., Usenix 2015) is based on efficient OT extension. We observe that our OPRF can be used to remove their PSI protocol's dependence on the bit-length of the parties' items. We implemented both PSI protocol variants and found ours to be 3.1–3.6× faster than Pinkas et al. for PSI of 128-bit strings and sufficiently large sets. Concretely, ours requires only 3.8 seconds to securely compute the intersection of 2^{20} -size sets, regardless of the bit length of the items. For very large sets, our protocol is only 4.3× slower than the *insecure* naïve hashing approach for PSI.

1 Introduction

This work involves OT, OPRF and PSI constructions. We start by reviewing the three primitives.


Oblivious Transfer. Oblivious Transfer (OT) has been a central primitive in the area of secure computation. Indeed, the original protocols of Yao [Yao86] and GMW [Gol04, GMW87] both use OT in a critical manner. In fact, OT is both necessary and sufficient for secure computation [Kil88]. Until early 2000's, the area of generic secure computation was often seen mainly as a feasibility exercise, and improving OT performance was not a priority research direction. This changed when Yao's Garbled Circuit (GC) was first implemented [MNPS04] and a surprisingly fast OT protocol (which we will call IKNP) was devised by Ishai et al. [IKNP03].

The IKNP OT extension protocol [IKNP03] is truly a gem; it allows 1-out-of-2 OT execution at the cost of computing and sending only a few hash values (but a security parameter of public key primitives evaluations were needed to bootstrap the system). IKNP was immediately noticed and since then universally used in implementations of the Yao and GMW protocols. It took a few years to realize that OT extension's use goes far beyond these fundamental applications. Many aspects of secure computation were strengthened and sped up by using OT extension. For example, Nielsen et al. [NNOB12] propose an approach to malicious two-party secure computation, which relates outputs and inputs of OTs in a larger construction. They critically rely on the low cost of batched OTs. Another example is the application of information-theoretic Gate Evaluation Secret Sharing (GESS) [Kol05] to the computational setting [KK12]. The idea of [KK12] is to stem the high cost in secret sizes of the GESS scheme by evaluating the circuit by shallow slices, and using OT extension to efficiently “glue” them together. Particularly relevant for our work, efficient OTs

*Bell Labs, kolesnikov@research.bell-labs.com

†MIT, rranjit@gmail.com

‡Oregon State University, {rosulekm, trieun}@eecs.oregonstate.edu



A

C

D

F

bins




B

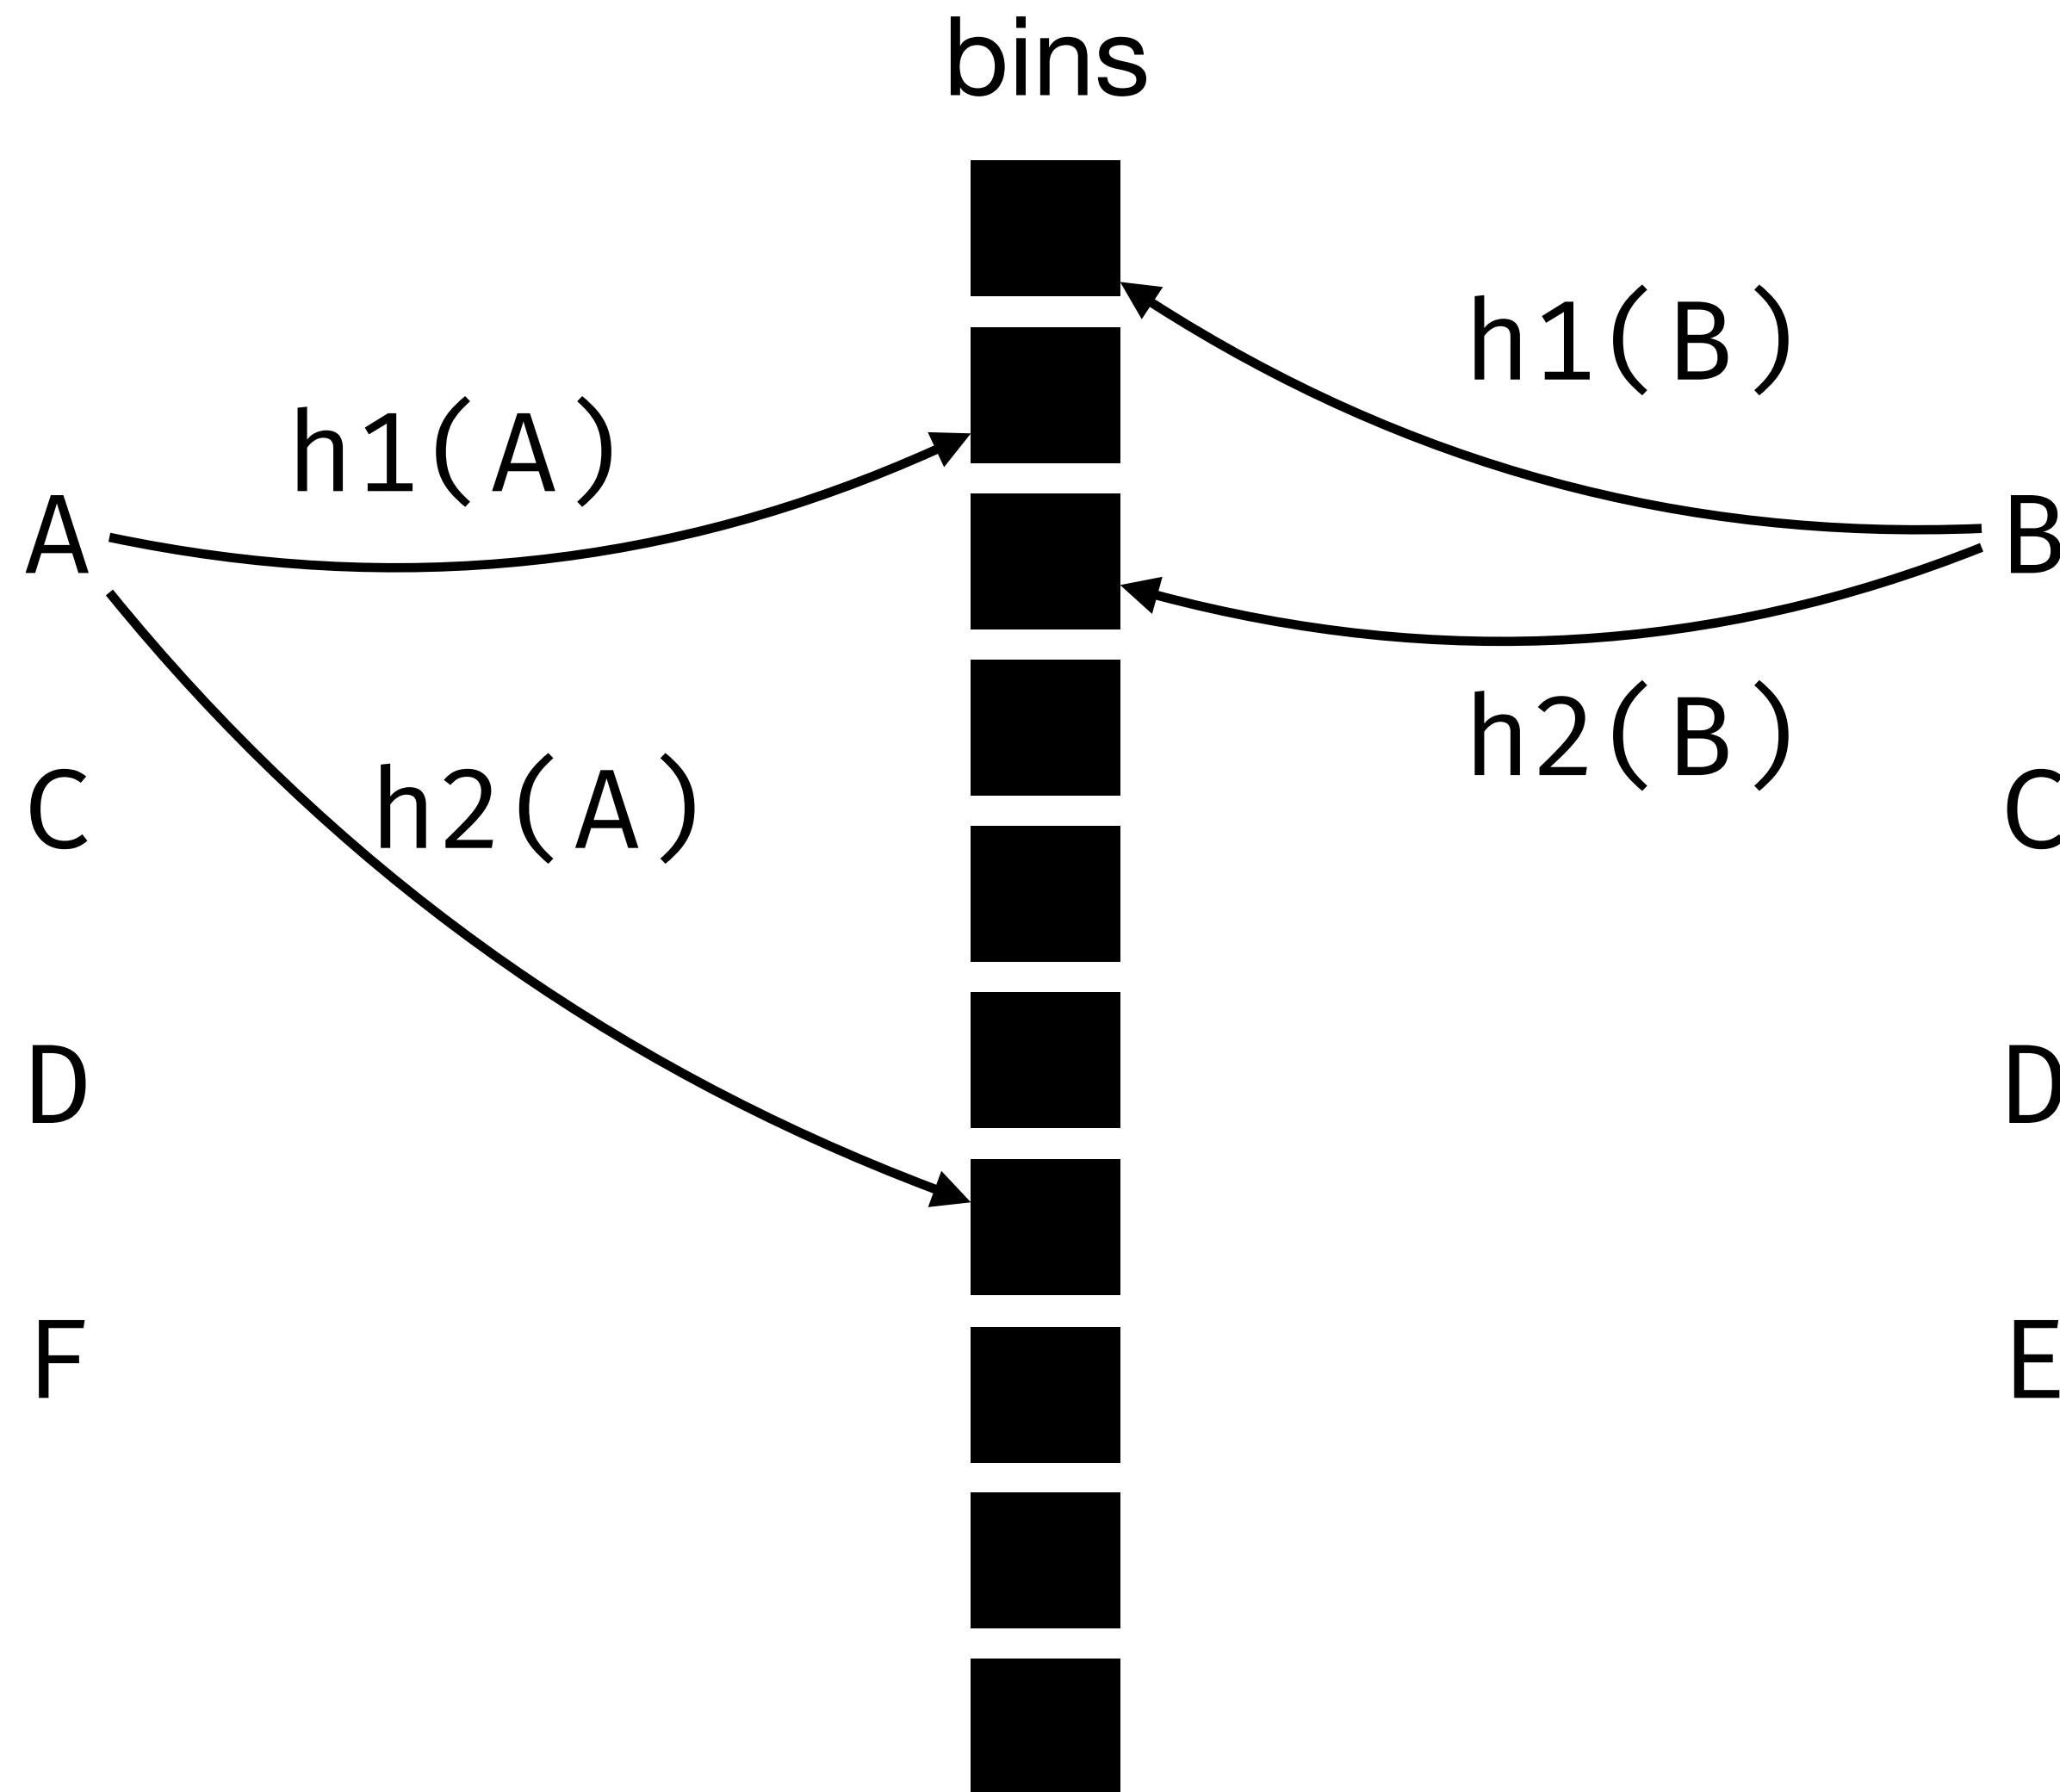

C

D



E


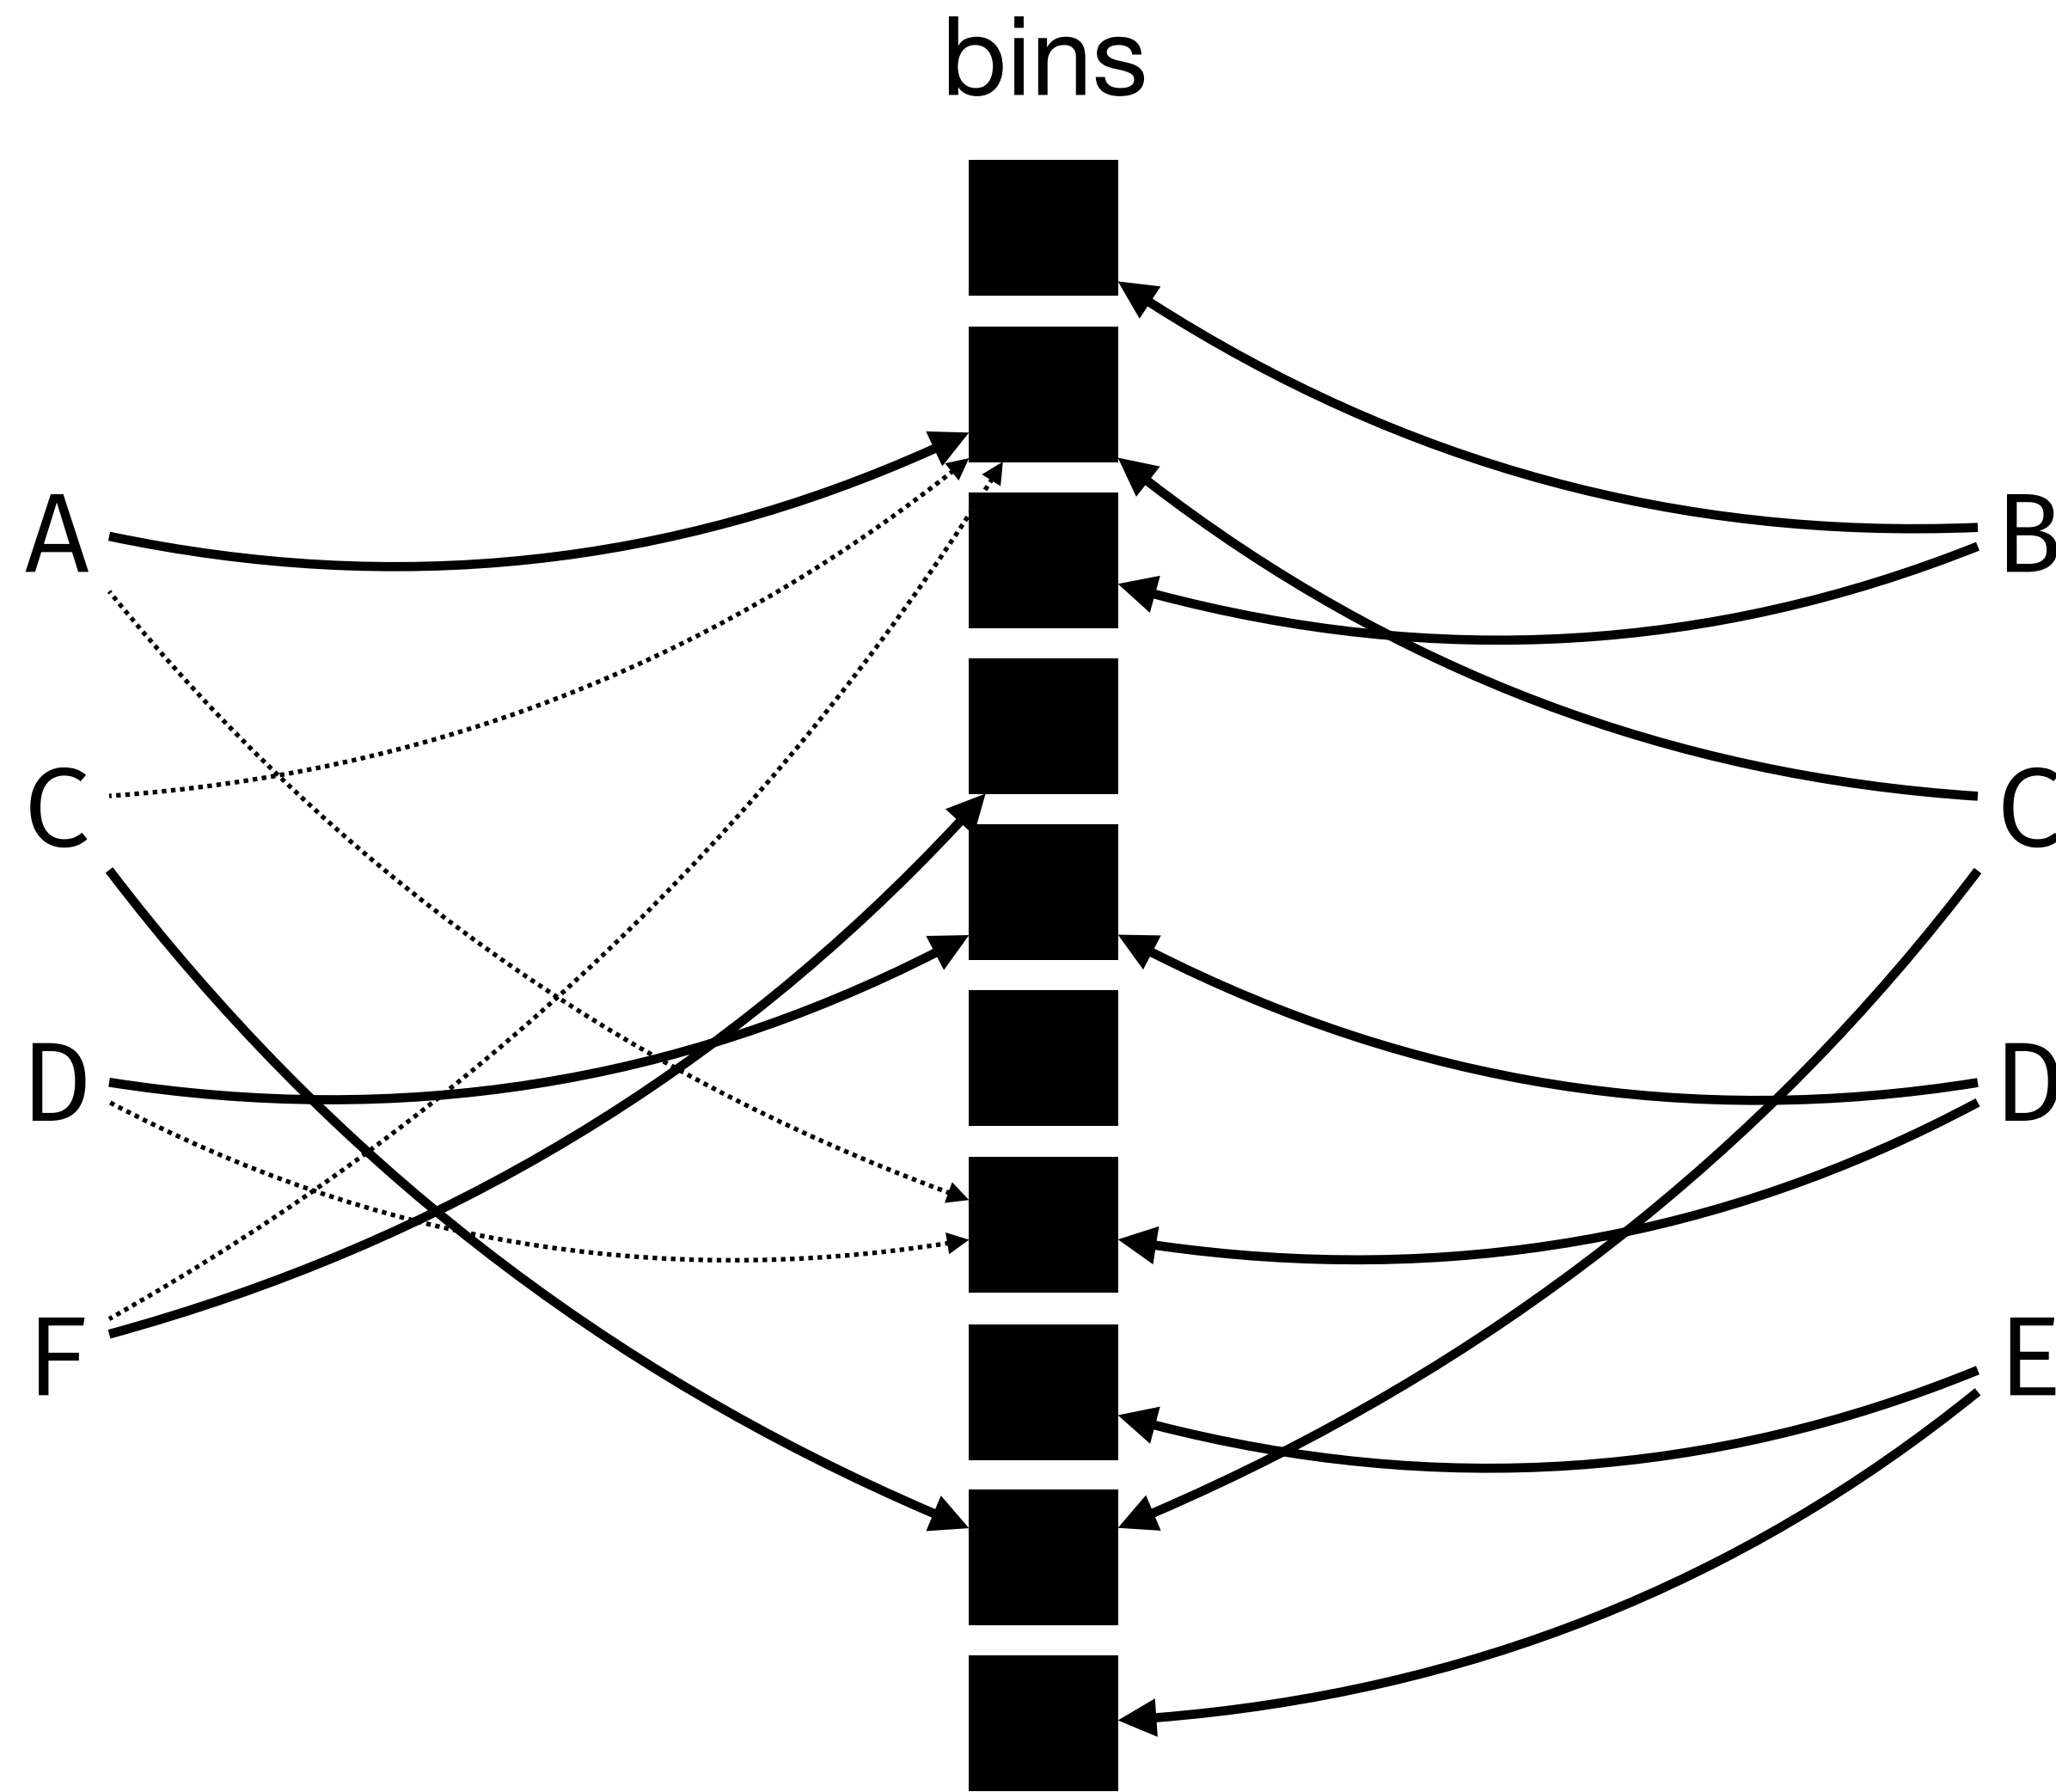


h1, h2 are two public hash functions





$h1, h2$ are two public hash functions









cuckoo hashing: if there are enough hash functions and enough bins, Alice can place at most one item in each bin (with constant overflow)



cuckoo hashing: if there are enough hash functions and enough bins, Alice can place at most one item in each bin (with very small overflow)





$F(k_0, B)$ $F(k_2, B)$ $F(k_1, C)$ **$F(k_8, C)$** $F(k_4, D)$...



Today's objectives

Define the private set intersection (PSI) problem

Construct PSI from MPC and discuss circuit PSI

Define Oblivious PRF

Construct PSI for large sets