Applications and Semi-Honest Security

CS 589 DH

Today's objectives

See applications of multiparty computation (MPC)

Sketch definition of ***semi-honest security***

Introduce the notion of a ***simulator***

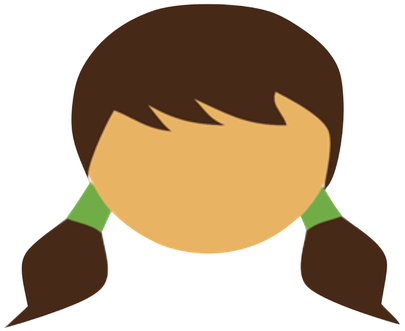

x

f

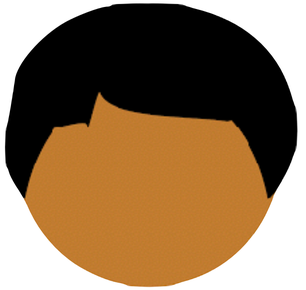

y

Crypto Magic

$f(x, y)$

$f(x, y)$

Privacy

Authenticity

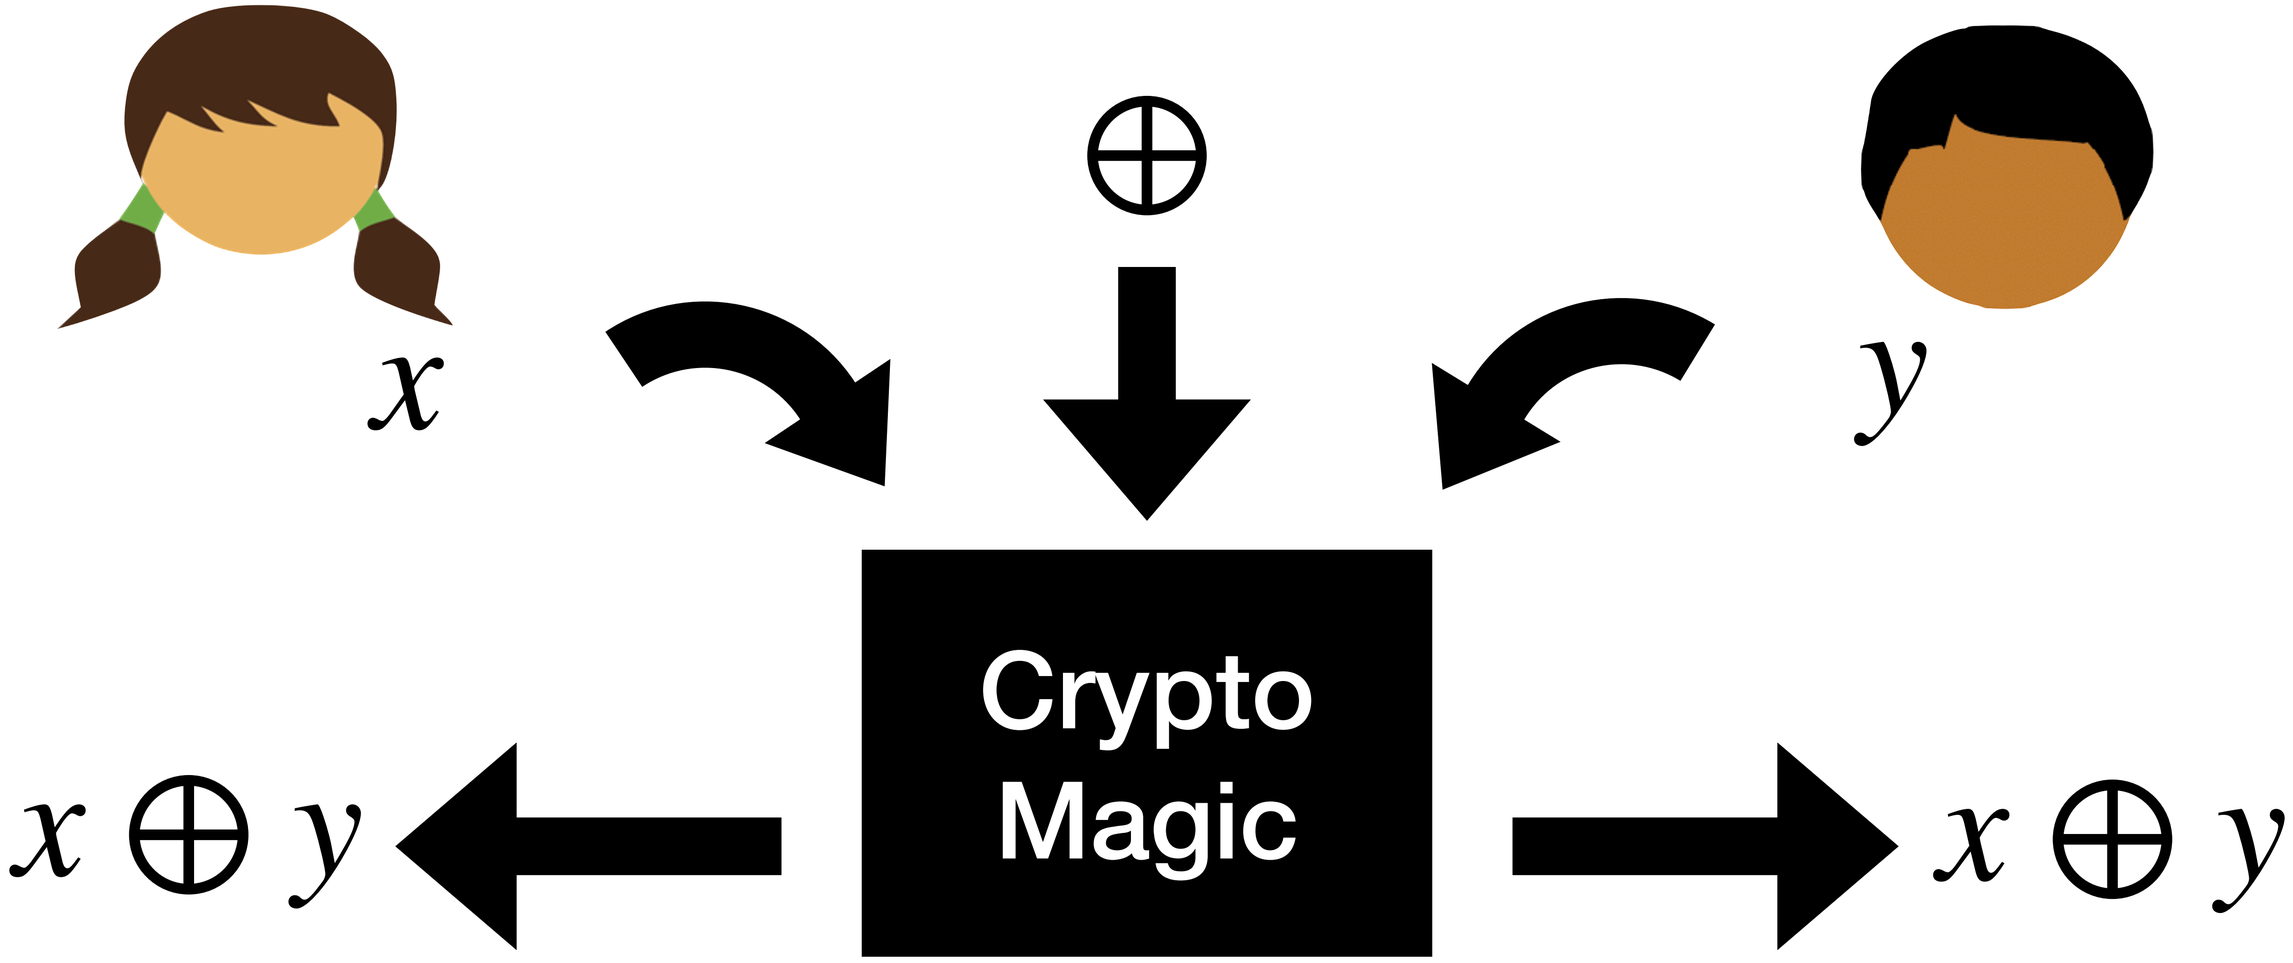
x
$\oplus$
y
Crypto
Magic
$x \oplus y$
$x \oplus y$

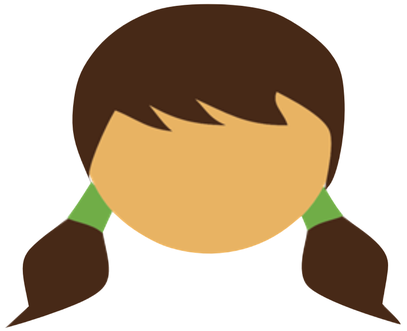

x

x

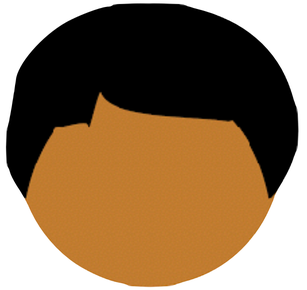

y

y

$x \oplus y$

$x \oplus y$

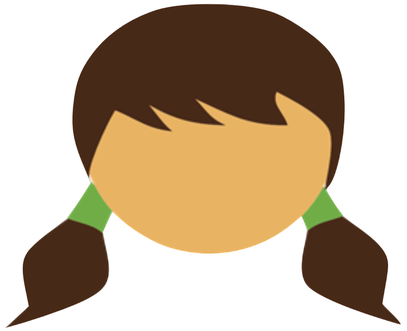

x

x

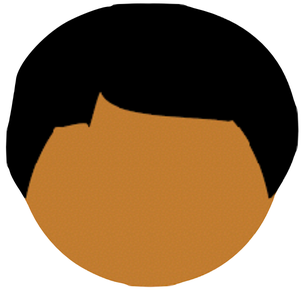

y

y

$x \oplus y$

$x \oplus y$

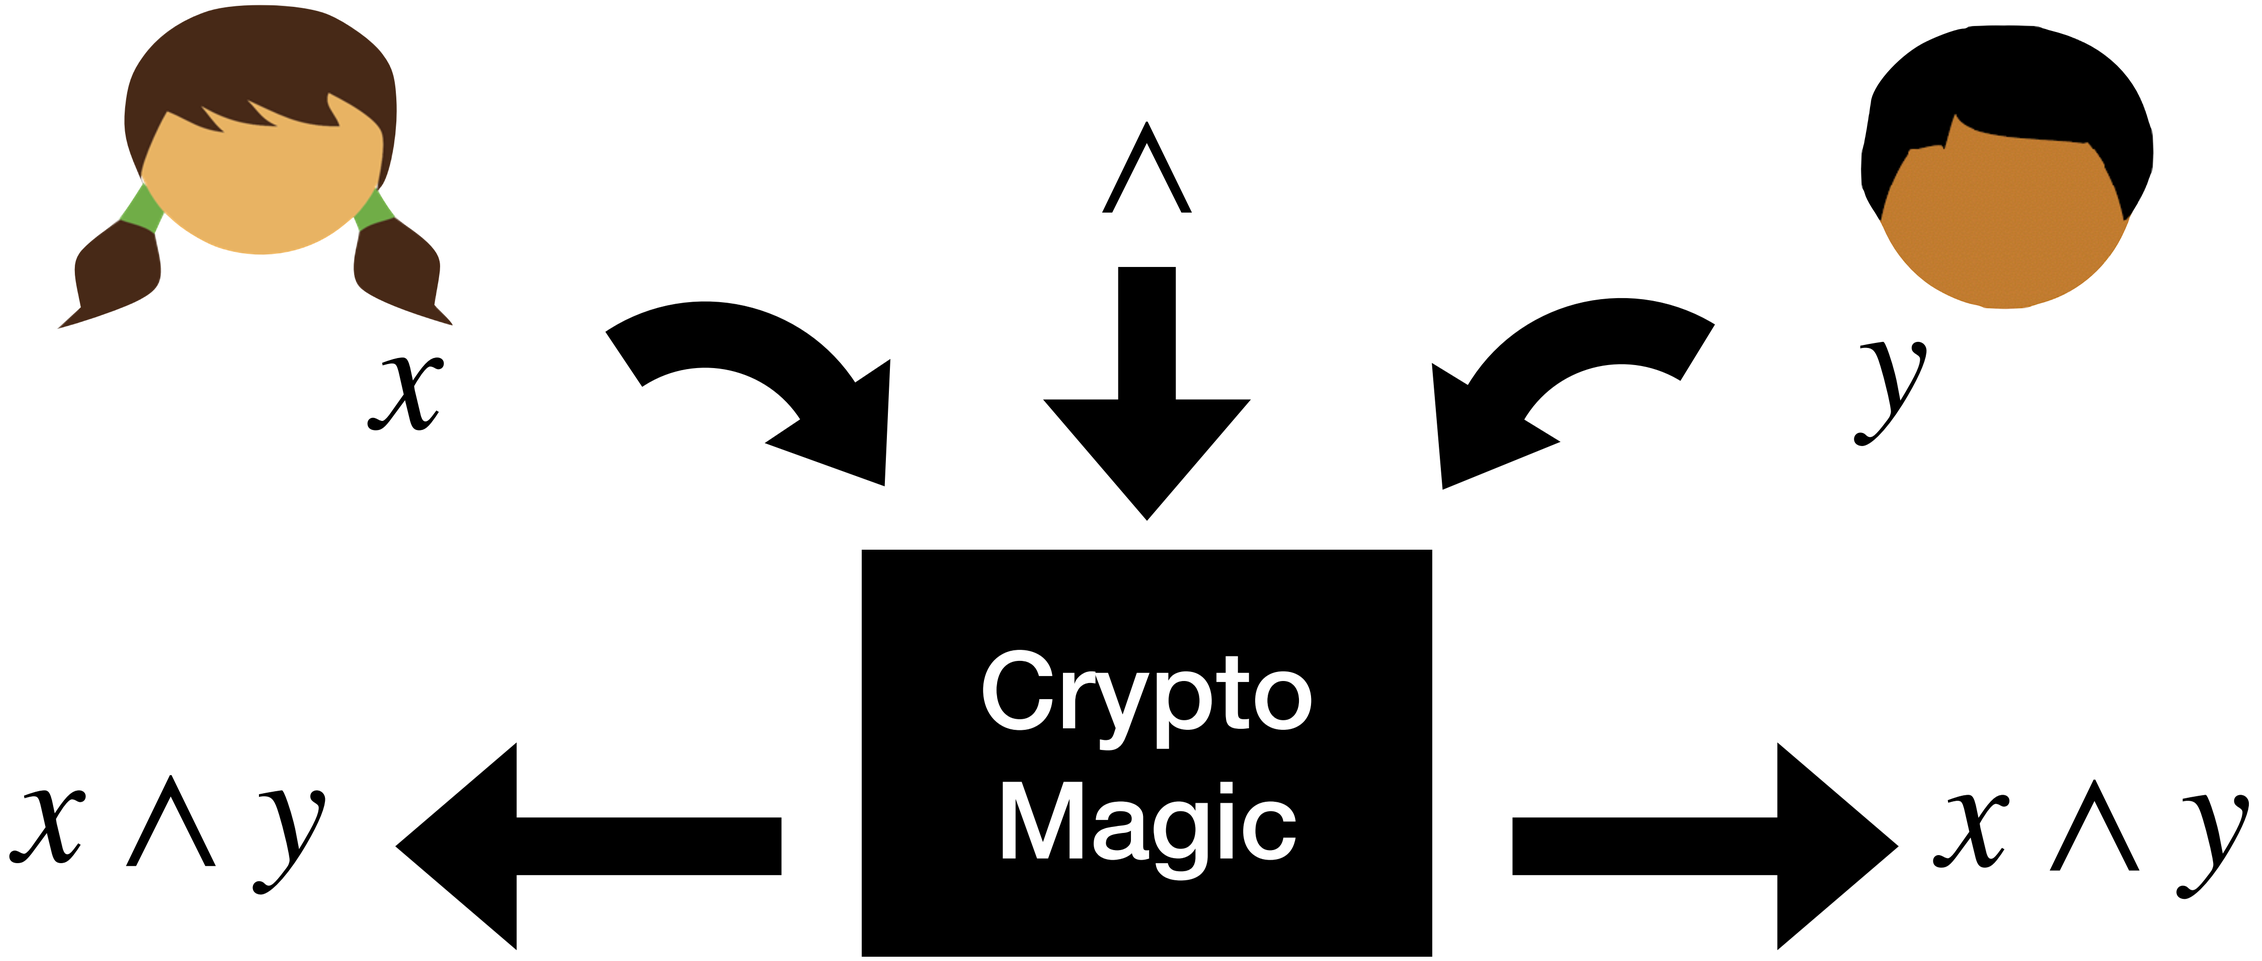
x
$\wedge$
y
Crypto
Magic
$x \wedge y$
$x \wedge y$

More Efficient Match-Making and Satisfiability

The Five Card Trick

Bert den Boer

Centrum voor Wiskunde en Informatica

Kruislaan 413, 1098 SJ Amsterdam, The Netherlands.

Abstract. A two-party cryptographic protocol for evaluating any binary gate is presented. It is more efficient than previous two-party computations, and can even perform single-party (i.e. satisfiability) proofs more efficiently than known techniques. As in all earlier multiparty computations and satisfiability protocols, commitments are a fundamental building block. Each party in our approach encodes a single input bit as 2 bit commitments. These are then combined to form 5 bit commitments, which are permuted, and can then be opened to reveal the output of the gate.

A Matchmaking Example

Alice and Bob never met before and wish to find out whether they have some particular mutual interest. But naturally each refuses to show interest first, because of the risk of getting an embarrassing "no" from the other. More formally, Alice has a secret bit a and Bob has a secret bit b and a protocol is needed that reveals exactly the logical "AND" of the two bits. Consequently, if Bob's bit is zero he should learn nothing about Alice's bit; if his bit is one, he cannot fail to learn Alice's bit because in that case her bit has the same value as the AND.

One way to achieve the desired protocol is by physical means—more precisely, five cards. The back of all cards are , as usual, the same. The face side of two of the cards are identical, say the two-of-hearts, and the face of the other three are identical, say the two-of-spades.

Initially,each party is given one card of each type and the remaining spade is put face down on the table and Bob then puts his cards face down on top of the initial spade. His secret choice of ordering for the two cards encodes his bit b: heart on top means 1 and the other way round means 0.

J.J. Quisquater and J. Vandewalle (Eds.): Advances in Cryptology - EUROCRYPT '89, LNCS 434, pp. 208-217, 1990.

The Five-Card Trick Can Be Done with Four Cards

Takaaki Mizuki, Michihito Kumamoto, and Hideaki Sone

Cyberscience Center, Tohoku University,
Aramaki-Aza-Aoba 6-3, Aoba-ku, Sendai 980-8578, Japan
tm-paper+ac2012@g-mail.tohoku-university.jp

Abstract. The "five-card trick" invented by Boer allows Alice and Bob to securely compute the AND function of their secret inputs using five cards—three black cards and two red cards—with identical backs. This paper shows that such a secure computation can be done with only four cards. Specifically, we give a protocol to achieve a secure computation of AND using only four cards—two black and two red. Our protocol is optimal in the sense that the number of required cards is minimum.

1 Introduction

Assume that two *honest-but-curious* players Alice and Bob, who hold secret bits $a \in \{0,1\}$ and $b \in \{0,1\}$, respectively, wish to *securely compute* the AND function, that is, they want to learn the value of $a \wedge b$ without revealing more of their own secret bits than necessary. The "five-card trick" invented in 1989 by Boer [2] achieves such a secure computation of AND using five cards $\clubsuit\clubsuit\clubsuit\heartsuit\heartsuit$. Now, after over two decades since the invention of the five-card trick, this paper improves upon the result: we show that the same secure computation can be done using only four cards $\clubsuit\clubsuit\heartsuit\heartsuit$.

This paper begins with an overview of the five-card trick.

1.1 The five-card trick

The "five-card trick" by Boer [2] is an elegant secure AND computation protocol that uses three $\clubsuit$s and two $\heartsuit$s. Before going into the details of the protocol, we first mention the properties of cards appearing in this paper.

All cards of the same type ($\clubsuit$ or $\heartsuit$) are assumed to be indistinguishable from one another. We use $?$ to denote a card lying face down. We also assume that the back $?$ of each card is identical. To deal with Boolean values, we use the following encoding:

$$\clubsuit\heartsuit = 0, \quad \heartsuit\clubsuit = 1. \tag{1}$$

Given a bit $x \in \{0,1\}$, a pair of face-down cards $??$ whose value is equal to x (according to the encoding rule (1) above) is called a *commitment to* x, and

MPC Applications

Secure Auctions

Secure Multiparty Computation Goes Live*

Peter Bogetoft[†], Dan Lund Christensen[¶], Ivan Damgård[‡], Martin Geisler[‡], Thomas Jakobsen[¶], Mikkel Krøigaard [‡], Jaus Dam Nielsen[‡], Jesper Buus Nielsen[‡], Kurt Nielsen[‡], Jakob Pagter[¶], Michael Schwartzbach[‡], and Tomas Toft[††]

[†] Inst. of Food and Resource Economics, University of Copenhagen
[‡] Department of Computer Science, University of Aarhus
§Dept. of Economics, Copenhagen Business School
¶The Alexandra Institute
†† CWI Amsterdam and TU/e

Abstract. In this note, we report on the first large-scale and practical application of multiparty computation, which took place in January 2008. We also report on the novel cryptographic protocols that were used.

1 Introduction and History

In multiparty computation (MPC), we consider a number of players $P_1, \ldots, P_n$, who initially each hold inputs $x_1, \ldots, x_n$, and we then want to securely compute some function f on these inputs, where $f(x_1, \ldots, x_n) = (y_1, \ldots, y_n)$, such that P_i learns y_i but no other information. This should hold, even if players exhibit some amount of adversarial behavior. The goal can be accomplished by an interactive protocol π that the players execute. Intuitively, we want that executing π is equivalent to having a trusted party $\mathcal{T}$ that receives privately x_i from P_i, computes the function, and returns y_i to each P_i [1]. With such a protocol we can - in principle - solve virtually any cryptographic protocol problem. The general theory of MPC was founded in the late 80-ties [16, 3, 7]. The theory was later developed in several ways – see for instance [21, 18, 8]. An overview of the theoretical results known can be found in [6].

Despite the obvious potential that MPC has in solving a wide range of problems, we have seen virtually no practical applications of MPC in the past. This is probably in part due to the fact that direct implementation of the first general protocols would lead to very inefficient solutions. Another factor has been a general lack of understanding in the general public of the potential of the technology. A lot of research has gone into solving the efficiency problems, both for general protocols [11, 17, 9] and for special types of computations such as voting [4, 12].

A different line of research has had explicit focus on a range of economic applications, which are particularly interesting for practical use. This approach was taken, for instance, by two research projects that the authors of this paper have been involved in: SCET (Secure Computing, Economy and Trust)[2] and SIMAP (Secure Information Management and Processing),[3], which has been responsible for the practical application of MPC described in this paper. In the economic field of mechanism design the concept of a trusted third party has been a central assumption since the 70's [15, 19, 10]. Ever since the field was initiated it has grown in momentum and turned into a truly cross disciplinary field. Today, many practical mechanisms require a trusted third party and it is natural to consider the possibility of implementing such a party using MPC. In particular, we have considered:

- Various types of auctions that involves sealed bids for different reasons. The most well-known is probably the standard highest bid auction with sealed bids, however, in terms of turnover another common variant is the so called double auction with many sellers and buyers. This auction handles scenarios where one wants to find a fair market price for a commodity given the existing supply and demand in the market.

* This work was sponsored by the Danish Strategic Research Council.
[1] This "equivalence" can be formalized using, for instance, Canetti's Universal Composability framework[5].
[2] see http://sikkerhed.alexandra.dk/uk/projects/scet
[3] see http://sikkerhed.alexandra.dk/uk/projects/simap

Web-based Multi-Party Computation with Application to Anonymous Aggregate Compensation Analytics

Andrei Lapets Eric Dunton Kyle Holzinger Frederick Jansen Azer Bestavros

CS Dept., Boston University
111 Cummington Mall
Boston, MA USA 02215
{lapets, edunton, kholz, fjansen, best}@bu.edu

Abstract

We describe the definition, design, implementation, and deployment of a multi-party computation protocol and supporting web-based infrastructure. The protocol and infrastructure constitute a software application that allows groups of cooperating parties, such as companies or other organizations, to collect aggregate data for statistical analysis without revealing the data of individual participants. The application was developed specifically to support a Boston Women's Workforce Council (BWWC) study of the gender wage gap among employers within the Greater Boston Area. The application was deployed successfully to collect aggregate statistical data pertaining to compensation levels across genders and demographics at a number of participating organizations.

1 Introduction

Modern organizations, including companies, educational institutions, and governments agencies, have been collecting and analyzing data pertaining to their internal operations for some time and to great effect, such as in evaluating performance or improving efficiency. While this data is of great value to the organizations themselves, it is likely that novel insights valuable to multiple organizations, to policymakers, or to society at large can be derived by combining data from these multiple organizations and analyzing it as a single corpus.

Unfortunately, the data collected by organizations internally is often proprietary and confidential, and its release may be potentially deleterious to their interests. Furthermore, while organizations may have the option of releasing sensitive data selectively to specific agents entrusted with its analysis, this presents a security risk: how will the data be physically transferred in a secure way, how will it be housed during the analysis, and how will it be destroyed after an analysis is complete?

Secure multi-party computation (MPC) techniques have been known for decades at least as theoretical constructs [25], and recent efforts [19, 13, 16, 21, 23] are finally bringing us closer to a point at which these techniques will be available to end-users (i.e., organizations interested in collectively analyzing their sensitive data).

In this report, we describe the definition, design, implementation, and deployment of a multi-party computation protocol and supporting web-based infrastructure for analyzing compensation data (broken down by gender and demographics) from a collection of employer organizations. The secure multi-party computation protocol utilized for this application is of relatively modest

Secure Auctions

Privacy-preserving studies

Students and Taxes: a Privacy-Preserving Social Study Using Secure Computation

Dan Bogdanov[1], Liina Kamm[1], Baldur Kubo[1], Reimo Rebane[1], Ville Sokk[1] and Riivo Talviste[1,2]

[1] Cybernetica, Tartu, Estonia
{dan.bogdanov, liina.kamm, baldur.kubo, reimo.rebane, ville.sokk, riivo.talviste}@cyber.ee
[2] University of Tartu, Institute of Computer Science, Tartu, Estonia

Abstract. We describe the use of secure multi-party computation for performing a large-scale privacy-preserving statistical study on real government data. In 2015, statisticians from the Estonian Center of Applied Research (CentAR) conducted a big data study to look for correlations between working during university studies and failing to graduate in time. The study was conducted by linking the database of individual tax payments from the Estonian Tax and Customs Board and the database of higher education events from the Ministry of Education and Research. Data collection, preparation and analysis were conducted using the Sharemind secure multi-party computation system that provided end-to-end cryptographic protection to the analysis. Using ten million tax records and half a million education records in the analysis, this is the largest cryptographically private statistical study ever conducted on real data.

Keywords: privacy, statistics, secure multi-party computation, case study

1 **Introduction**

Information and communication technology (ICT) is a growing industry where highly skilled specialists are in demand. This causes concern to both industry, where the wages keep rising, and the academia that cannot often match the pay grades offered by the industry. The universities in Estonia formed a hypothesis that students who work during their studies, do not graduate in the allotted time. Moreover, many students quit before graduation, thus, not acquiring the skills needed for building more complex ICT systems.

In this paper, we describe a big data study on Estonian government data that researches this topic and uses privacy-enhancing technologies to protect personal data. We collaborated with a team of social scientists who designed a statistical study that links tax and education records to determine the working habits of both ICT and non-ICT students. However, running the actual study would normally be impossible, as data protection and tax secrecy legislation

Secure Auctions

Privacy-preserving studies

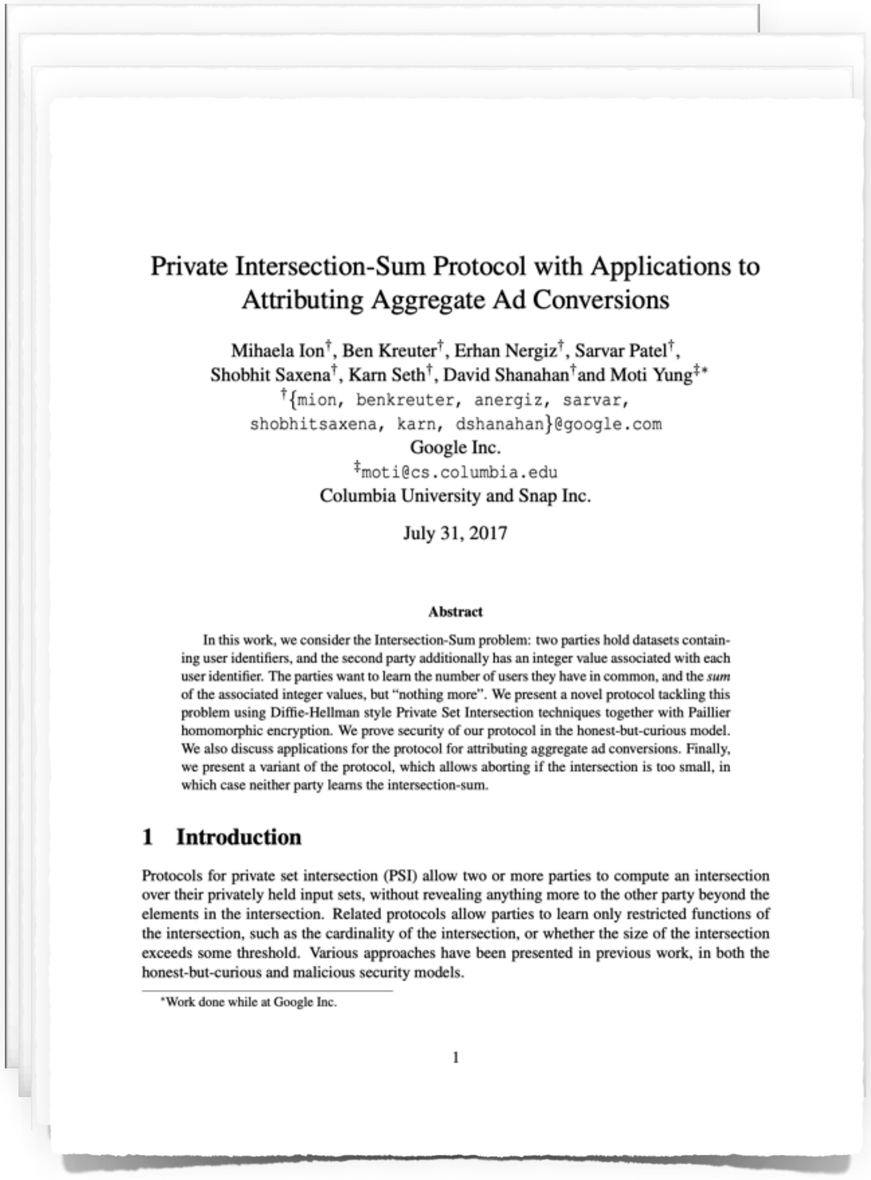

Private Intersection-Sum Protocol with Applications to Attributing Aggregate Ad Conversions

Mihaela Ion[†], Ben Kreuter[†], Erhan Nergiz[†], Sarvar Patel[†], Shobhit Saxena[†], Karn Seth[†], David Shanahan[†]and Moti Yung[‡*]
[†]{mion, benkreuter, anergiz, sarvar, shobhitsaxena, karn, dshanahan}@google.com
Google Inc.
[‡]moti@cs.columbia.edu
Columbia University and Snap Inc.

July 31, 2017

Abstract

In this work, we consider the Intersection-Sum problem: two parties hold datasets containing user identifiers, and the second party additionally has an integer value associated with each user identifier. The parties want to learn the number of users they have in common, and the *sum* of the associated integer values, but "nothing more". We present a novel protocol tackling this problem using Diffie-Hellman style Private Set Intersection techniques together with Paillier homomorphic encryption. We prove security of our protocol in the honest-but-curious model. We also discuss applications for the protocol for attributing aggregate ad conversions. Finally, we present a variant of the protocol, which allows aborting if the intersection is too small, in which case neither party learns the intersection-sum.

1 Introduction

Protocols for private set intersection (PSI) allow two or more parties to compute an intersection over their privately held input sets, without revealing anything more to the other party beyond the elements in the intersection. Related protocols allow parties to learn only restricted functions of the intersection, such as the cardinality of the intersection, or whether the size of the intersection exceeds some threshold. Various approaches have been presented in previous work, in both the honest-but-curious and malicious security models.

*Work done while at Google Inc.

1

Secure Auctions

Privacy-preserving studies

Privacy-preserving advertising

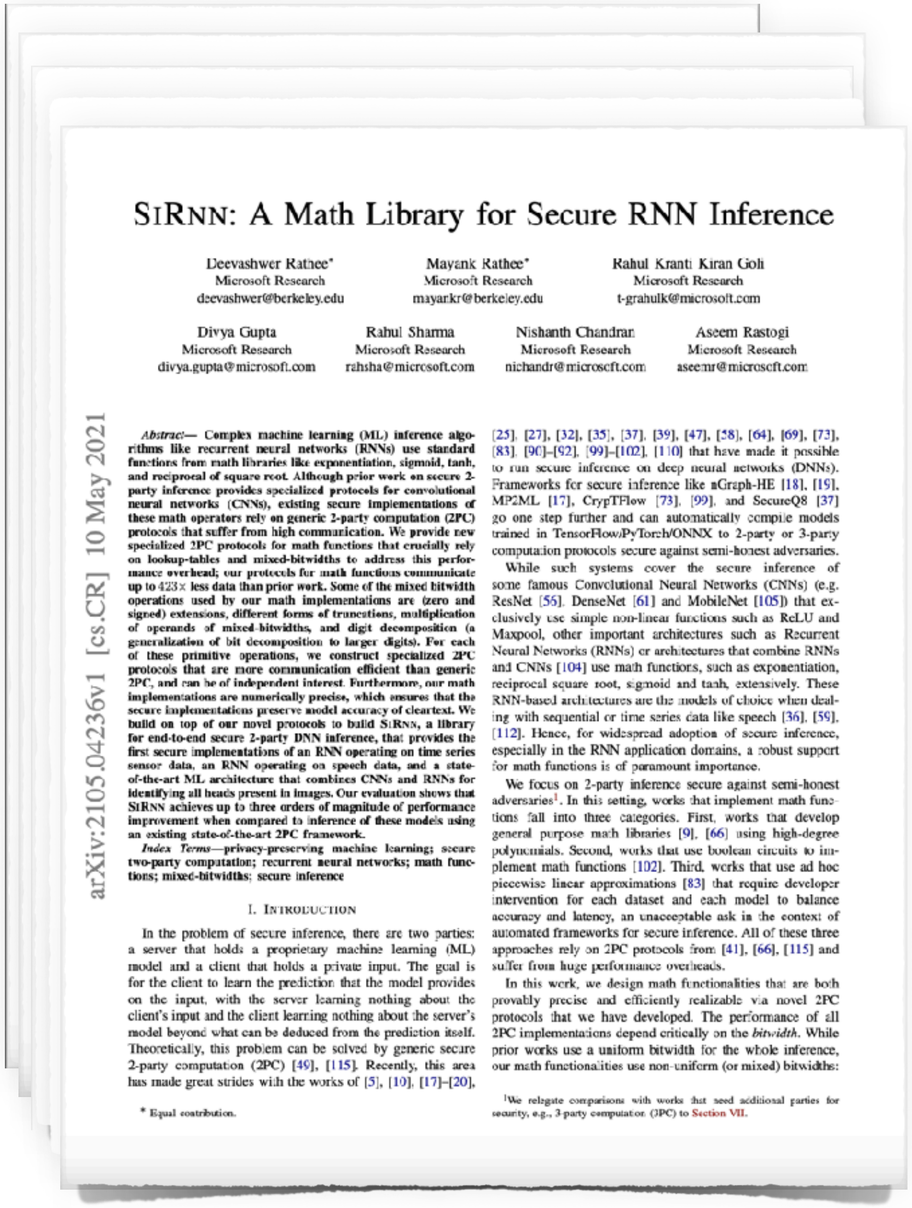

SIRNN: A Math Library for Secure RNN Inference

Deevashwer Rathee*
Microsoft Research
deevashwer@berkeley.edu

Mayank Rathee*
Microsoft Research
mayankr@berkeley.edu

Rahul Kranti Kiran Goli
Microsoft Research
t-grahulk@microsoft.com

Divya Gupta
Microsoft Research
divya.gupta@microsoft.com

Rahul Sharma
Microsoft Research
rahsha@microsoft.com

Nishanth Chandran
Microsoft Research
nichandr@microsoft.com

Aseem Rastogi
Microsoft Research
aseemr@microsoft.com

Abstract— **Complex machine learning (ML) inference algorithms like recurrent neural networks (RNNs) use standard functions from math libraries like exponentiation, sigmoid, tanh, and reciprocal of square root. Although prior work on secure 2-party inference provides specialized protocols for convolutional neural networks (CNNs), existing secure implementations of these math operators rely on generic 2-party computation (2PC) protocols that suffer from high communication. We provide new specialized 2PC protocols for math functions that crucially rely on lookup-tables and mixed-bitwidths to address this performance overhead; our protocols for math functions communicate up to 423× less data than prior work. Some of the mixed bitwidth operations used by our math implementations are (zero and signed) extensions, different forms of truncations, multiplication of operands of mixed-bitwidths, and digit decomposition (a generalization of bit decomposition to larger digits). For each of these primitive operations, we construct specialized 2PC protocols that are more communication efficient than generic 2PC, and can be of independent interest. Furthermore, our math implementations are numerically precise, which ensures that the secure implementations preserve model accuracy of cleartext. We build on top of our novel protocols to build SIRNN, a library for end-to-end secure 2-party DNN inference, that provides the first secure implementations of an RNN operating on time series sensor data, an RNN operating on speech data, and a state-of-the-art ML architecture that combines CNNs and RNNs for identifying all heads present in images. Our evaluation shows that SIRNN achieves up to three orders of magnitude of performance improvement when compared to inference of these models using an existing state-of-the-art 2PC framework.**

Index Terms—**privacy-preserving machine learning; secure two-party computation; recurrent neural networks; math functions; mixed-bitwidths; secure inference**

I. Introduction

In the problem of secure inference, there are two parties: a server that holds a proprietary machine learning (ML) model and a client that holds a private input. The goal is for the client to learn the prediction that the model provides on the input, with the server learning nothing about the client's input and the client learning nothing about the server's model beyond what can be deduced from the prediction itself. Theoretically, this problem can be solved by generic secure 2-party computation (2PC) [46], [115]. Recently, this area has made great strides with the works of [5], [10], [17]–[20], [25], [27], [32], [35], [37], [39], [47], [58], [64], [69], [73], [83], [90]–[92], [99]–[102], [110] that have made it possible to run secure inference on deep neural networks (DNNs). Frameworks for secure inference like nGraph-HE [18], [15], MP2ML [17], CrypTFlow [73], [99], and SecureQ8 [37] go one step further and can automatically compile models trained in TensorFlow/PyTorch/ONNX to 2-party or 3-party computation protocols secure against semi-honest adversaries.

While such systems cover the secure inference of some famous Convolutional Neural Networks (CNNs) (e.g. ResNet [55], DenseNet [61] and MobileNet [105]) that exclusively use simple non-linear functions such as ReLU and Maxpool, other important architectures such as Recurrent Neural Networks (RNNs) or architectures that combine RNNs and CNNs [104] use math functions, such as exponentiation, reciprocal square root, sigmoid and tanh, extensively. These RNN-based architectures are the models of choice when dealing with sequential or time series data like speech [36], [59], [112]. Hence, for widespread adoption of secure inference, especially in the RNN application domains, a robust support for math functions is of paramount importance.

We focus on 2-party inference secure against semi-honest adversaries[1]. In this setting, works that implement math functions fall into three categories. First, works that develop general purpose math libraries [9], [66] using high-degree polynomials. Second, works that use boolean circuits to implement math functions [102]. Third, works that use ad hoc piecewise linear approximations [83] that require developer intervention for each dataset and each model to balance accuracy and latency, an unacceptable ask in the context of automated frameworks for secure inference. All of these three approaches rely on 2PC protocols from [41], [66], [115] and suffer from huge performance overheads.

In this work, we design math functionalities that are both provably precise and efficiently realizable via novel 2PC protocols that we have developed. The performance of all 2PC implementations depend critically on the *bitwidth*. While prior works use a uniform bitwidth for the whole inference, our math functionalities use non-uniform (or mixed) bitwidths:

* Equal contribution.

[1]We relegate comparisons with works that need additional parties for security, e.g., 3-party computation (3PC) to Section VII.

Secure Auctions

Privacy-preserving studies

Privacy-preserving advertising

Privacy-preserving analytics

(Secure Machine Learning)

Differentially Private Secure Multi-Party Computation for Federated Learning in Financial Applications

David Byrd
db@gatech.edu
School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia

Antigoni Polychroniadou
antigoni.poly@jpmorgan.com
J.P. Morgan AI Research
New York, New York

ABSTRACT

Federated Learning enables a population of clients, working with a trusted server, to collaboratively learn a shared machine learning model while keeping each client's data within its own local systems. This reduces the risk of exposing sensitive data, but it is still possible to reverse engineer information about a client's private data set from communicated model parameters. Most federated learning systems therefore use differential privacy to introduce noise to the parameters. This adds uncertainty to any attempt to reveal private client data, but also reduces the accuracy of the shared model, limiting the useful scale of privacy-preserving noise. A system can further reduce the coordinating server's ability to recover private client information, without additional accuracy loss, by also including secure multiparty computation. An approach combining both techniques is especially relevant to financial firms as it allows new possibilities for collaborative learning without exposing sensitive client data. This could produce more accurate models for important tasks like optimal trade execution, credit origination, or fraud detection. The key contributions of this paper are: We present a privacy-preserving federated learning protocol to a non-specialist audience, demonstrate it using logistic regression on a real-world credit card fraud data set, and evaluate it using an open-source simulation platform which we have adapted for the development of federated learning systems.

KEYWORDS

federated learning, simulation, multiagent, finance, privacy

ACM Reference Format:
David Byrd and Antigoni Polychroniadou. 2020. Differentially Private Secure Multi-Party Computation for Federated Learning in Financial Applications. In *ACM International Conference on AI in Finance (ICAIF '20), October 15–16, 2020, New York, NY, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3383455.3422562

ACM ISBN 978-1-4503-7584-9/20/10…$15.00
https://doi.org/10.1145/3383455.3422562

1 INTRODUCTION

Modern financial firms routinely need to conduct analysis of large data sets stored across multiple servers or devices. A typical response is to combine these data sets into a single central database, but this approach introduces a number of privacy challenges: The institution may not have appropriate authority or permission to transfer locally stored information, the owner of the data may not *want* it shared, and centralization of the data may worsen the potential consequences of a data breach.

For example, the mobile app ai.type collected personal data from its users' phones and uploaded this information to a central database. Security researchers gained access to the database and obtained the names, email addresses, passwords, and other sensitive information of 31 million users of the Android version of the app. Such incidents highlight the risks and challenges associated with centralized data solutions. [5]

In this section, we motivate our approach while providing an extensive non-technical overview of the underlying techniques.

1.1 Federated Learning

One approach to mitigate the mentioned privacy concerns is to analyze the multiple data sets separately and share only the resulting insights from each analysis. This approach is realized in a recently-introduced technique called federated analysis. [2] Federated *learning*, already adopted by large companies like Google, allows users to share insights (perhaps the parameters of a trained model) from the data on their laptops or mobile devices *without* ever sharing the data itself, typically as follows:

1. Users train a local model on their individual data.
2. Each user sends their model weights to a trusted server.
3. The server computes an average-weight shared model.
4. The shared model is returned to all of the users.
5. Users retrain a local model starting from the shared model.

For instance, email providers could use federated learning to reduce the amount of spam their customers receive. Instead of each provider using its own spam filter trained from its customers' reported spam email, the providers could combine their models to create a shared spam-detection mechanism, without sharing their individual customers' reported spam emails. For a survey of recent advances in federated learning, see Kairouz et al. [13]

It is still possible, however, for a malicious party to potentially compromise the privacy of the individual users by inferring details of a training data set from the trained model's weights or parameters [16, 19]. It is important to protect sensitive user information while still providing highly accurate inferences.

Secure Auctions

Privacy-preserving studies

Privacy-preserving advertising

Privacy-preserving analytics

(Secure Machine Learning)

Financial Fraud Detection

Differentially Private Secure Multi-Party Computation for Federated Learning in Financial Applications

David Byrd
db@gatech.edu
School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia

Antigoni Polychroniadou
antigoni.poly@jpmorgan.com
J.P. Morgan AI Research
New York, New York

ABSTRACT

Federated Learning enables a population of clients, working with a trusted server, to collaboratively learn a shared machine learning model while keeping each client's data within its own local systems. This reduces the risk of exposing sensitive data, but it is still possible to reverse engineer information about a client's private data set from communicated model parameters. Most federated learning systems therefore use differential privacy to introduce noise to the parameters. This adds uncertainty to any attempt to reveal private client data, but also reduces the accuracy of the shared model, limiting the useful scale of privacy-preserving noise. A system can further reduce the coordinating server's ability to recover private client information, without additional accuracy loss, by also including secure multiparty computation. An approach combining both techniques is especially relevant to financial firms as it allows new possibilities for collaborative learning without exposing sensitive client data. This could produce more accurate models for important tasks like optimal trade execution, credit origination, or fraud detection. The key contributions of this paper are: We present a privacy-preserving federated learning protocol to a non-specialist audience, demonstrate it using logistic regression on a real-world credit card fraud data set, and evaluate it using an open-source simulation platform which we have adapted for the development of federated learning systems.

KEYWORDS

federated learning, simulation, multiagent, finance, privacy

ACM Reference Format:
David Byrd and Antigoni Polychroniadou. 2020. Differentially Private Secure Multi-Party Computation for Federated Learning in Financial Applications. In *ACM International Conference on AI in Finance (ICAIF '20), October 15–16, 2020, New York, NY, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3383455.3422562

ACM ISBN 978-1-4503-7584-9/20/10…$15.00
https://doi.org/10.1145/3383455.3422562

1 INTRODUCTION

Modern financial firms routinely need to conduct analysis of large data sets stored across multiple servers or devices. A typical response is to combine these data sets into a single central database, but this approach introduces a number of privacy challenges: The institution may not have appropriate authority or permission to transfer locally stored information, the owner of the data may not *want* it shared, and centralization of the data may worsen the potential consequences of a data breach.

For example, the mobile app ai.type collected personal data from its users' phones and uploaded this information to a central database. Security researchers gained access to the database and obtained the names, email addresses, passwords, and other sensitive information of 31 million users of the Android version of the app. Such incidents highlight the risks and challenges associated with centralized data solutions. [5]

In this section, we motivate our approach while providing an extensive non-technical overview of the underlying techniques.

1.1 Federated Learning

One approach to mitigate the mentioned privacy concerns is to analyze the multiple data sets separately and share only the resulting insights from each analysis. This approach is realized in a recently-introduced technique called federated analysis. [2] Federated *learning*, already adopted by large companies like Google, allows users to share insights (perhaps the parameters of a trained model) from the data on their laptops or mobile devices *without* ever sharing the data itself, typically as follows:

1. Users train a local model on their individual data.
2. Each user sends their model weights to a trusted server.
3. The server computes an average-weight shared model.
4. The shared model is returned to all of the users.
5. Users retrain a local model starting from the shared model.

For instance, email providers could use federated learning to reduce the amount of spam their customers receive. Instead of each provider using its own spam filter trained from its customers' reported spam email, the providers could combine their models to create a shared spam-detection mechanism, without sharing their individual customers' reported spam emails. For a survey of recent advances in federated learning, see Kairouz et al. [13]

It is still possible, however, for a malicious party to potentially compromise the privacy of the individual users by inferring details of a training data set from the trained model's weights or parameters [16, 19]. It is important to protect sensitive user information while still providing highly accurate inferences.

Secure Auctions

Privacy-preserving studies

Privacy-preserving advertising

Privacy-preserving analytics

(Secure Machine Learning)

Financial Fraud Detection

…and much more

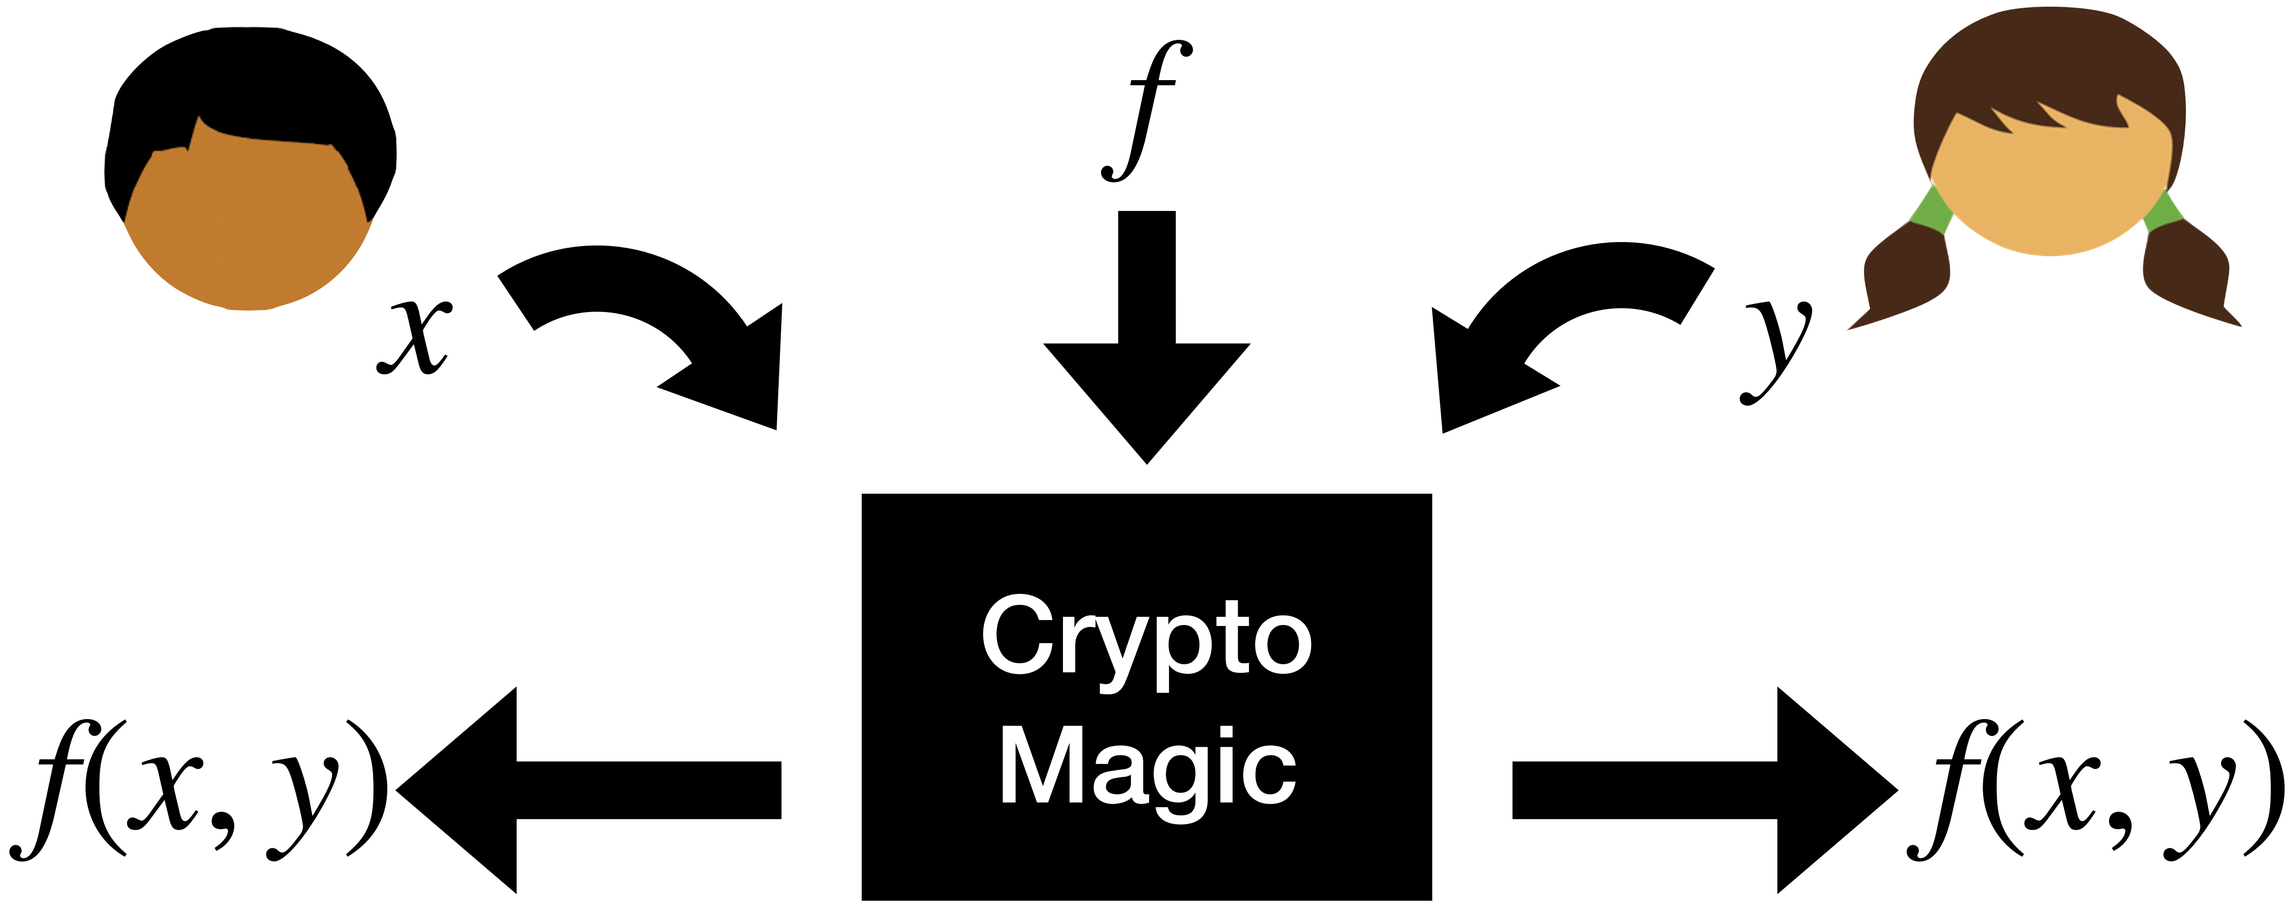

Privacy
Authenticity

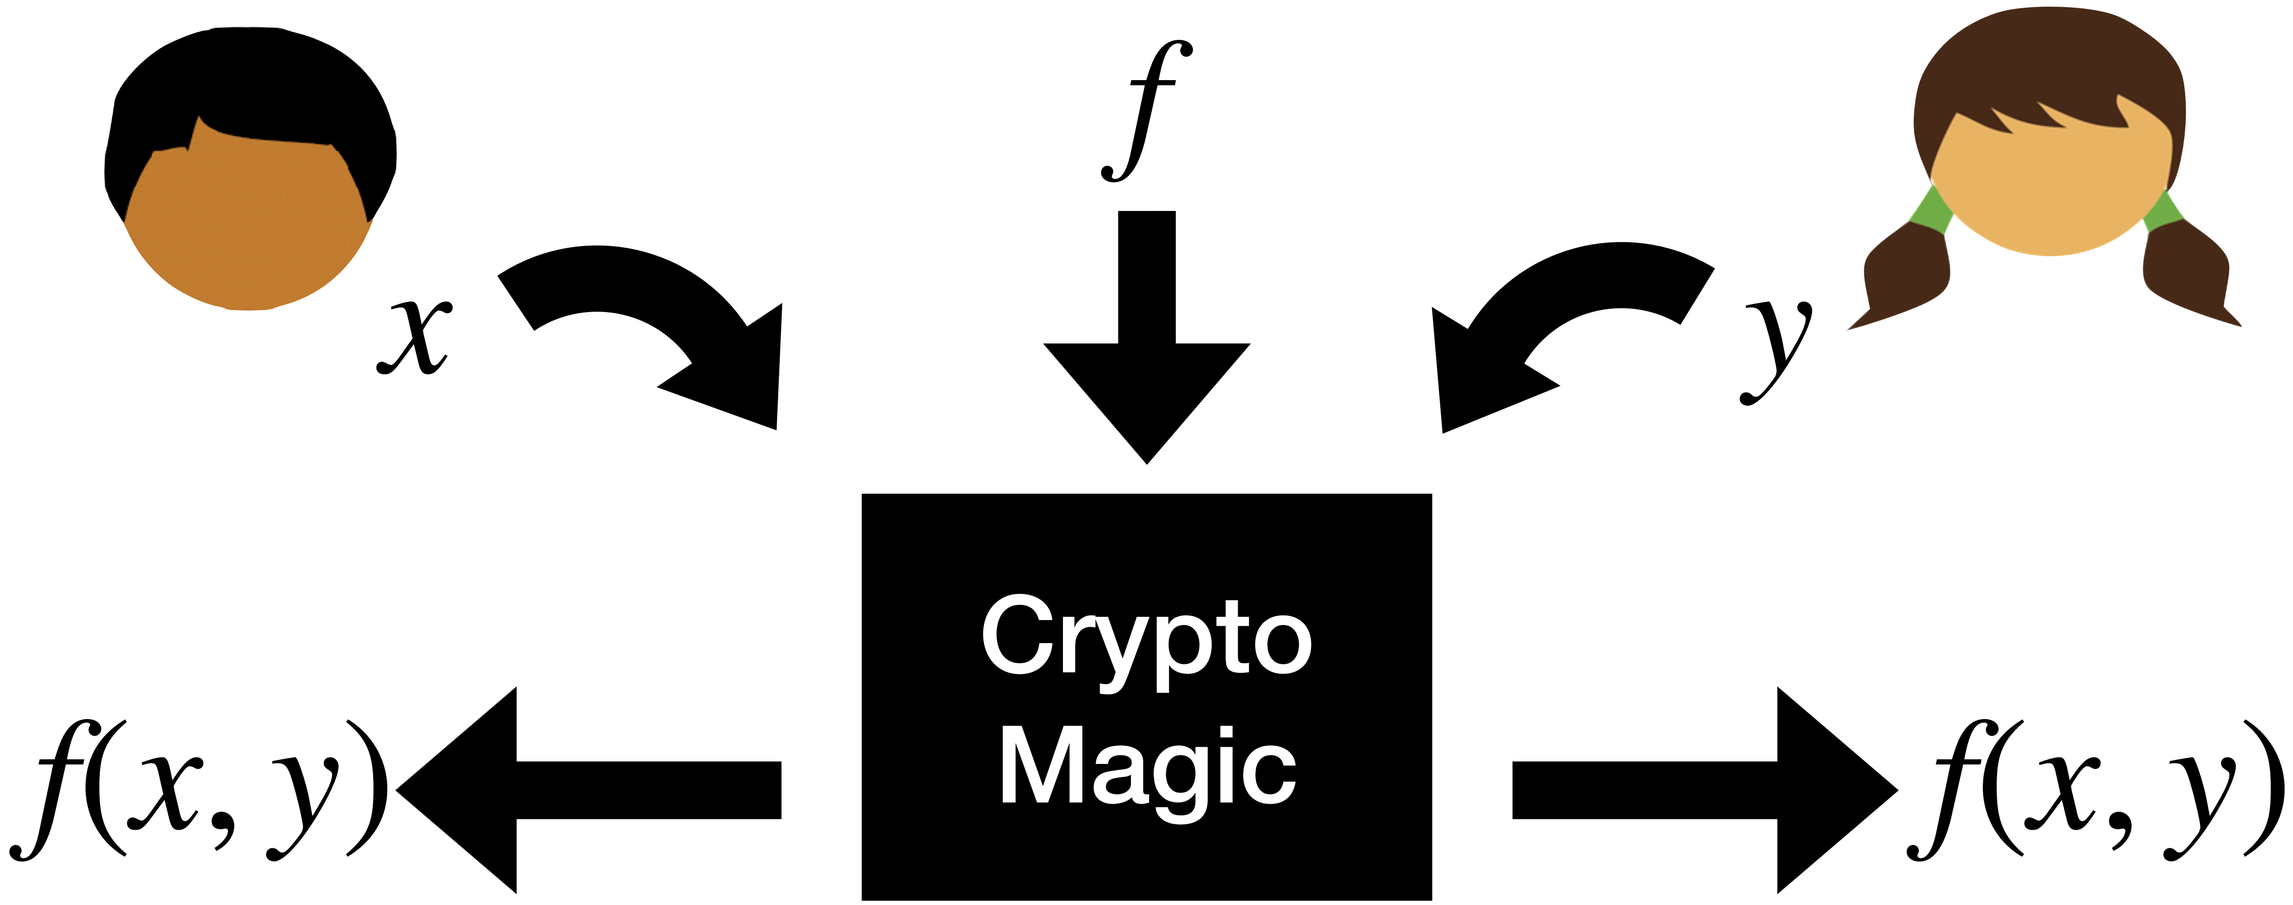

Privacy

~~Authenticity~~

Privacy
~~Authenticity~~

Semi-honest
Passive
Honest-but-curious

Privacy +
Authenticity

Malicious
Active

What is a Protocol

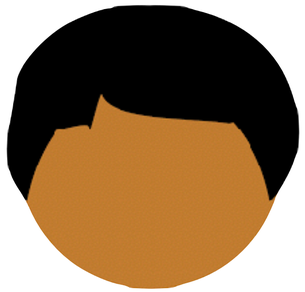

We will model parties as
interactive, *randomized algorithms*

$r \in_\$ \{0,1\}^*$

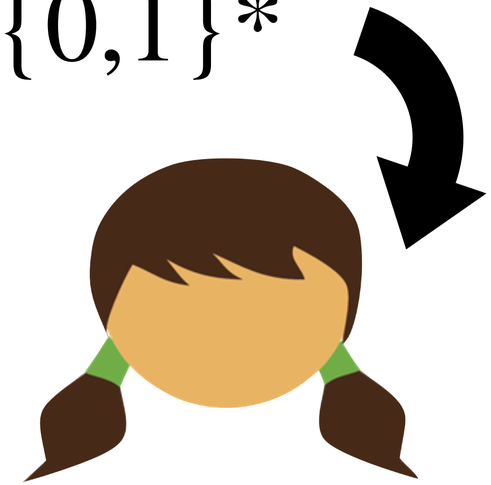

We will model parties as
interactive, *randomized algorithms*

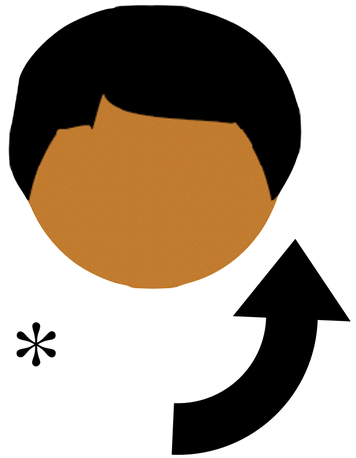

$s \in_\$ \{0,1\}^*$

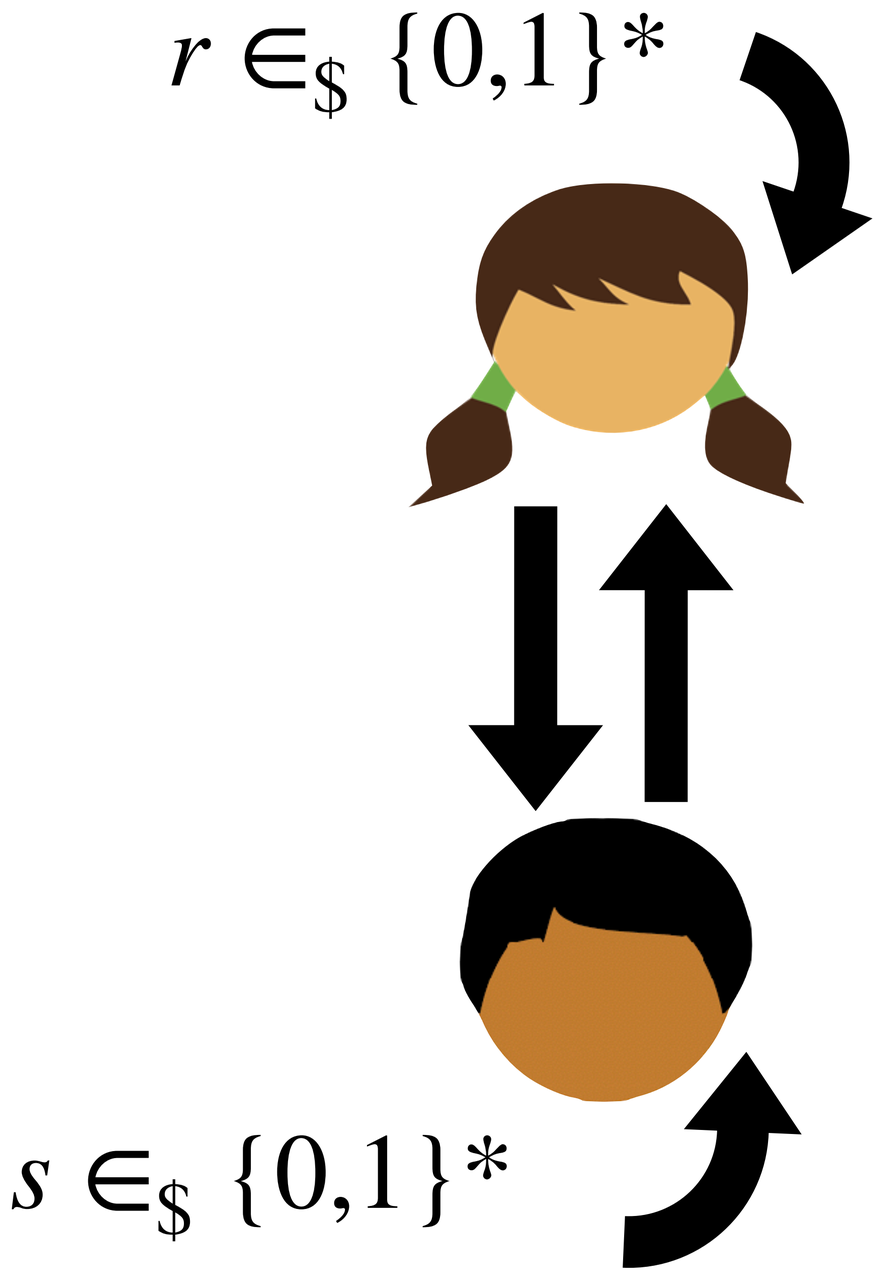

We will model parties as *interactive*, *randomized algorithms*

Unless otherwise stated, we will assume *secure point-to-point channels between parties* and that the network is synchronous

Π

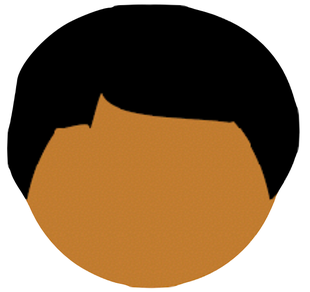

Π

x, r

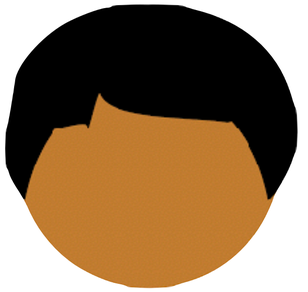

y, s

Π

x, r

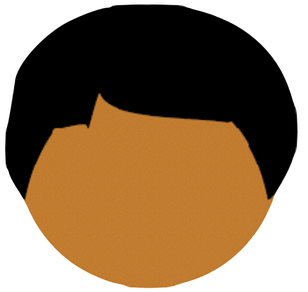

y, s

Π

x, r

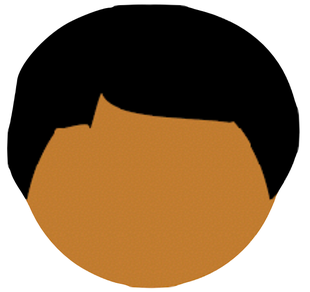

y, s

Π

x, r

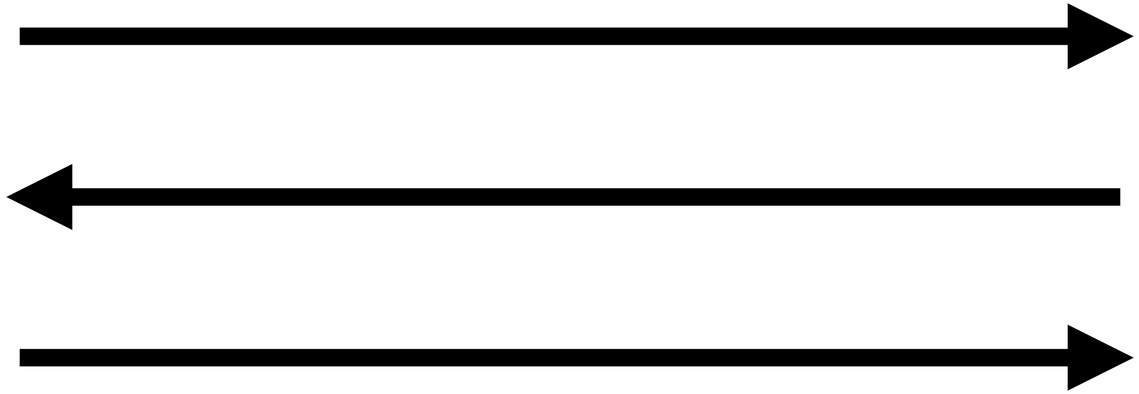

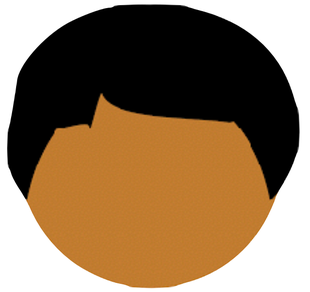

y, s

Π

x, r

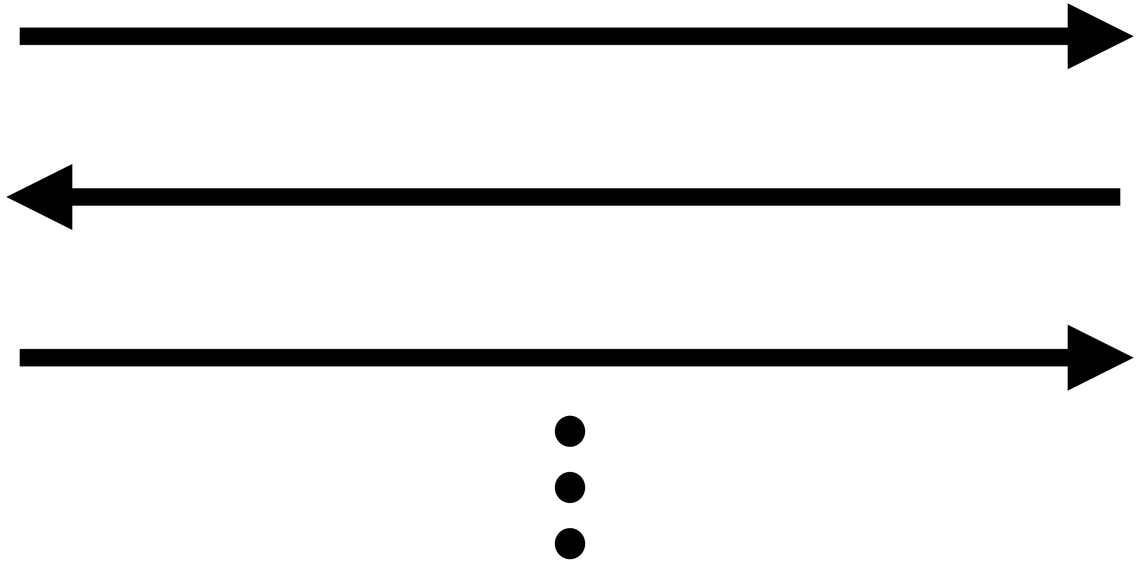

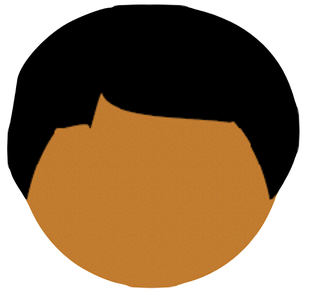

y, s

$output_A$

$output_B$

Π

x, r

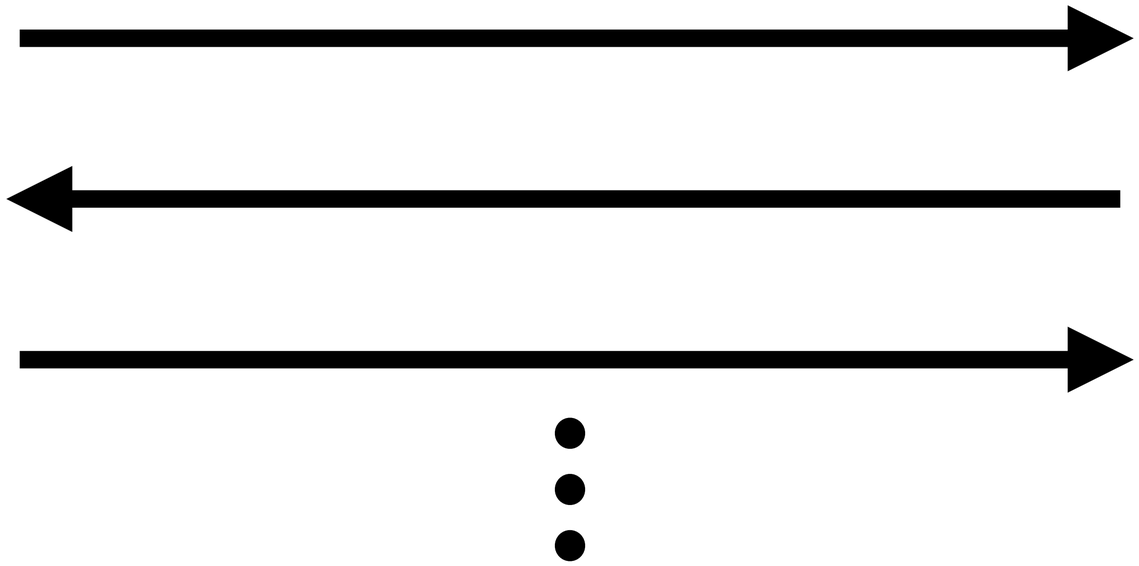

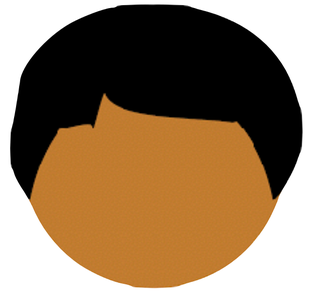

y, s

$$output_A = f(x, y)$$

$$output_B = g(x, y)$$

Correctness: if parties follow the protocol, they get the correct output

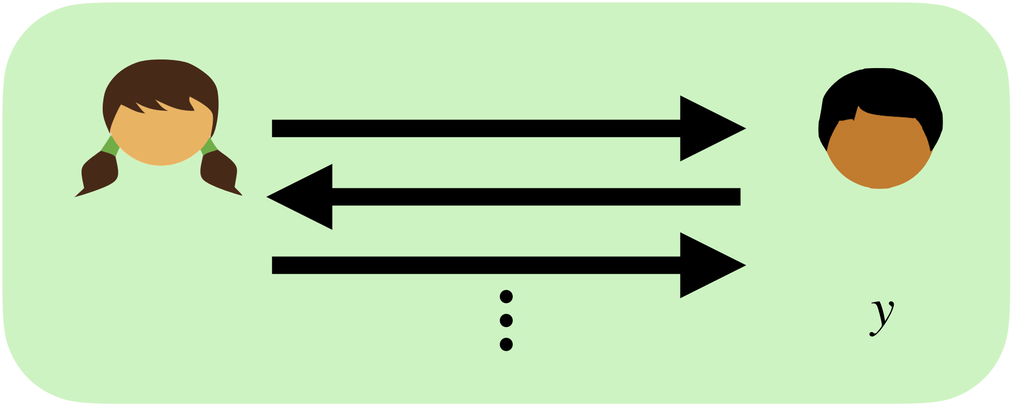

Privacy (informal)

Semi-honest Security

Bob learns nothing about Alice's input, except for what is implied by the function output (and vice versa)

Π

x, s

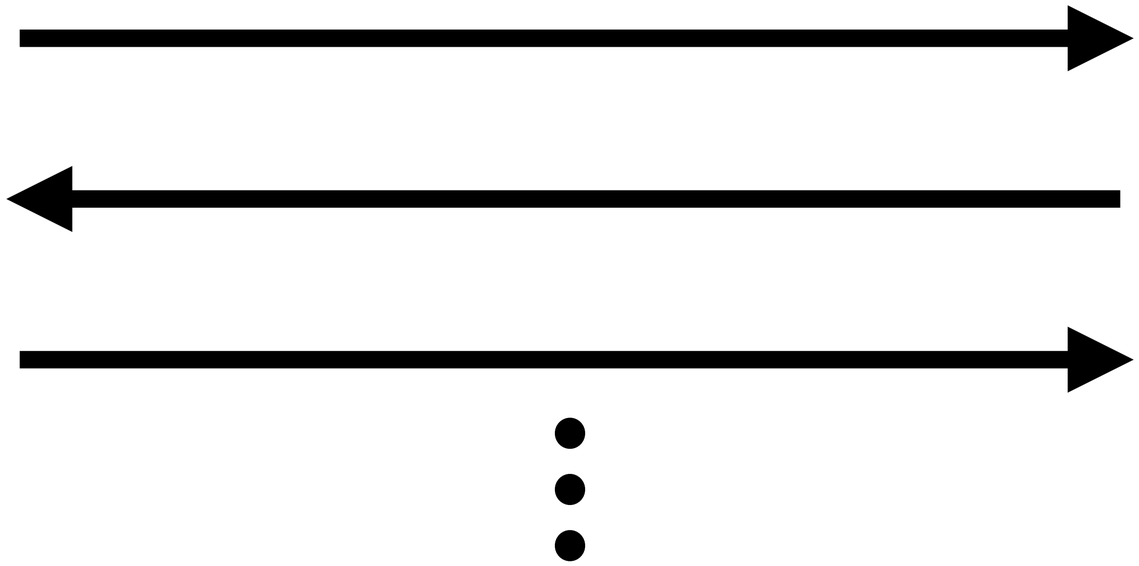

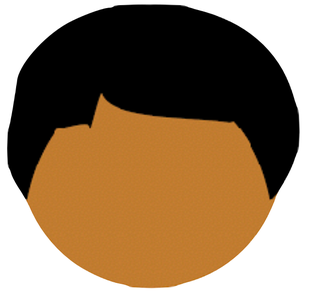

y, r

Π

x

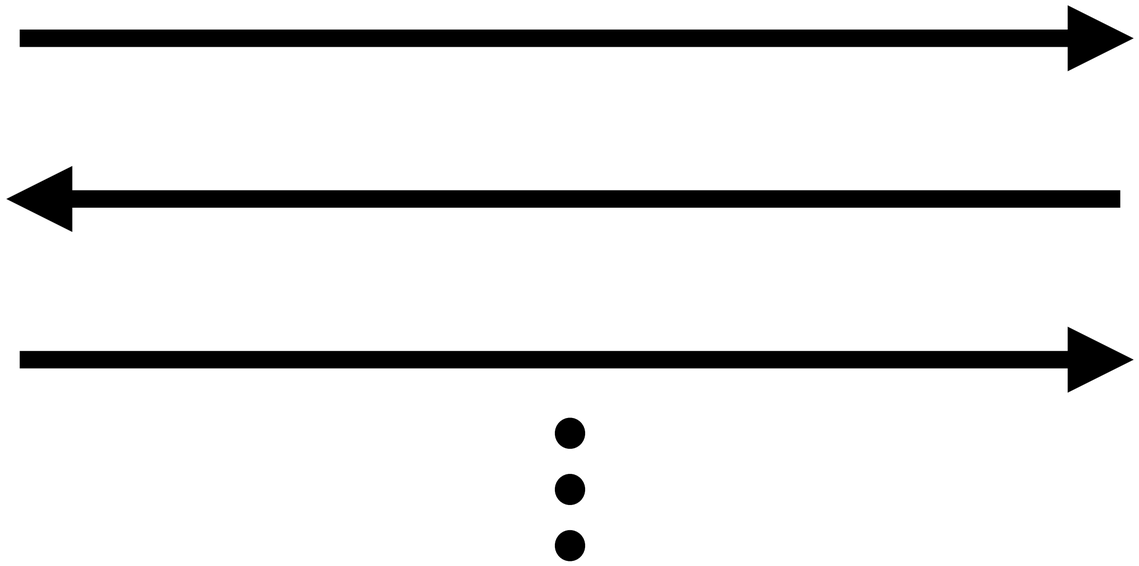

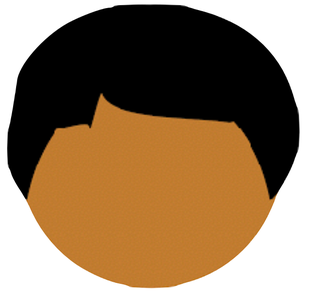

y

Π

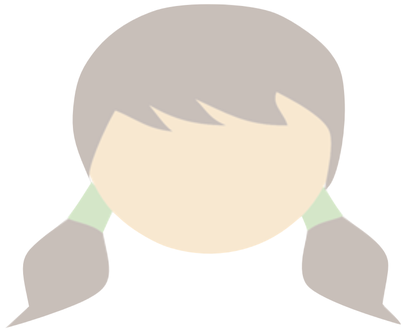

x

m_0

m_1

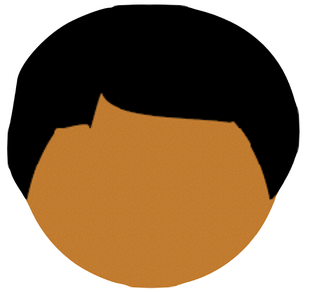

y

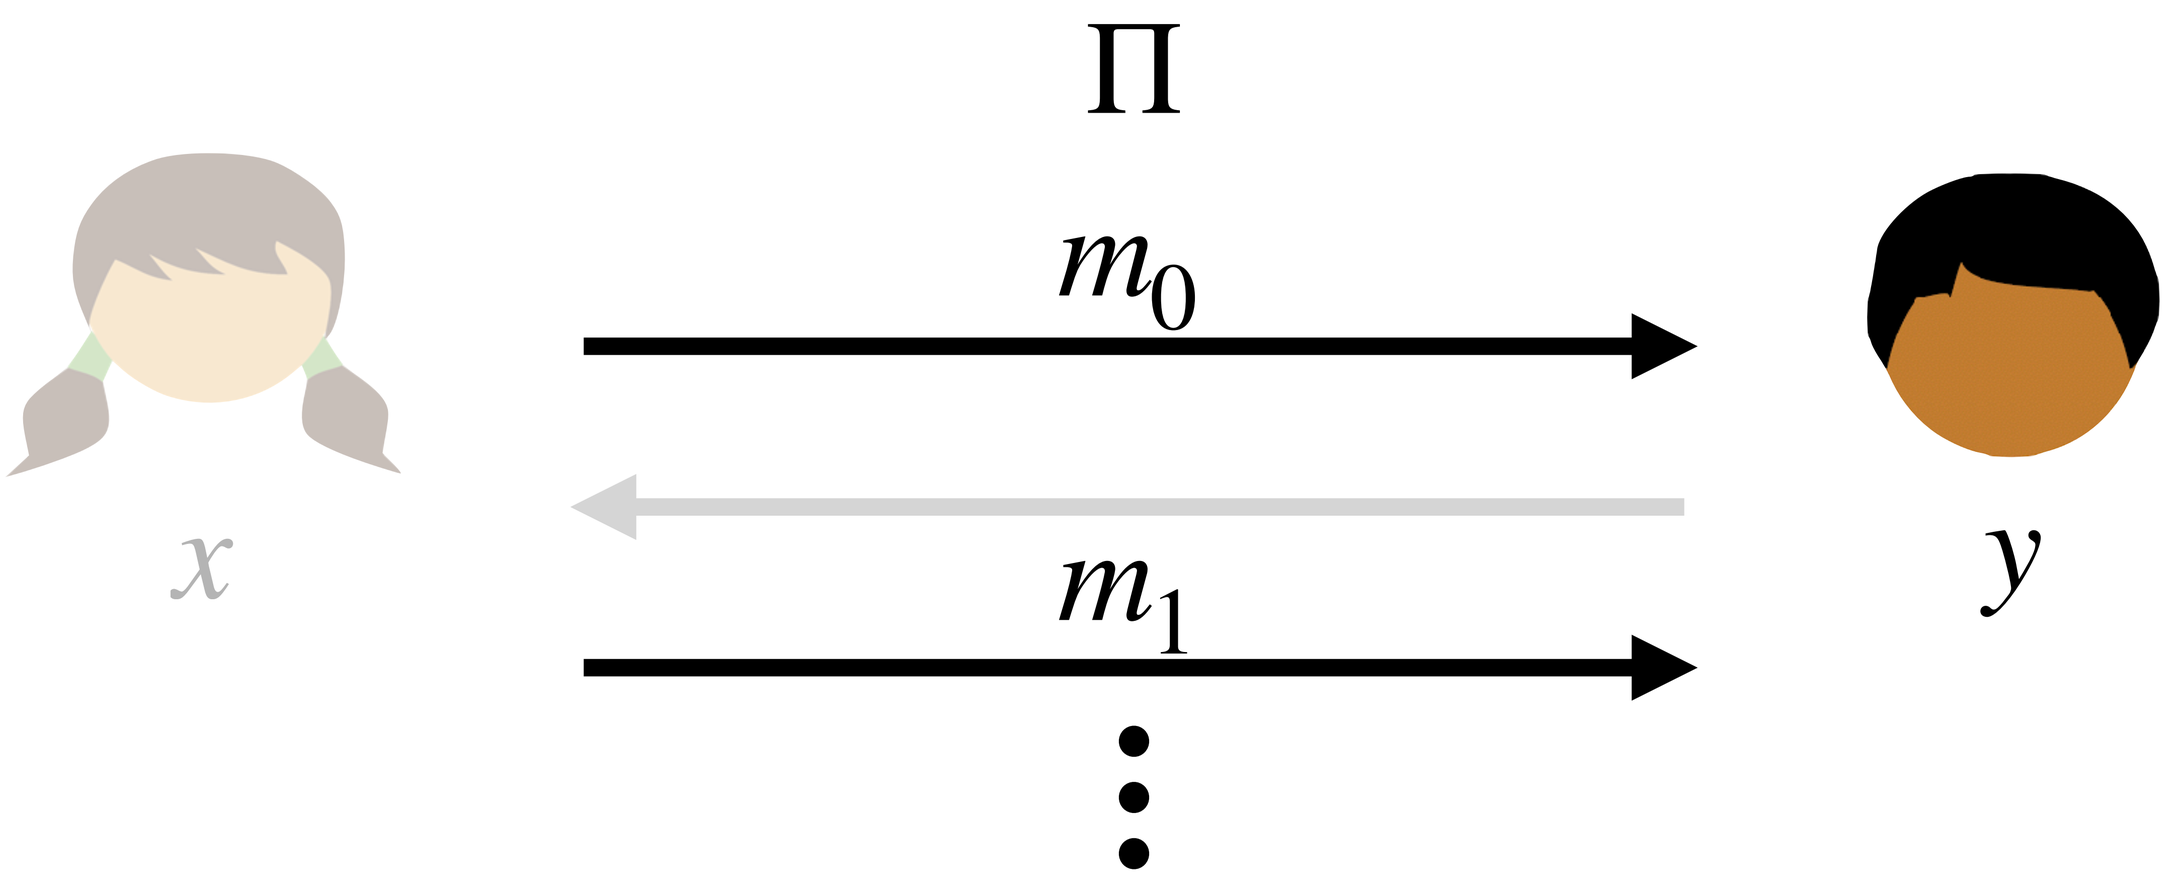

$$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$$

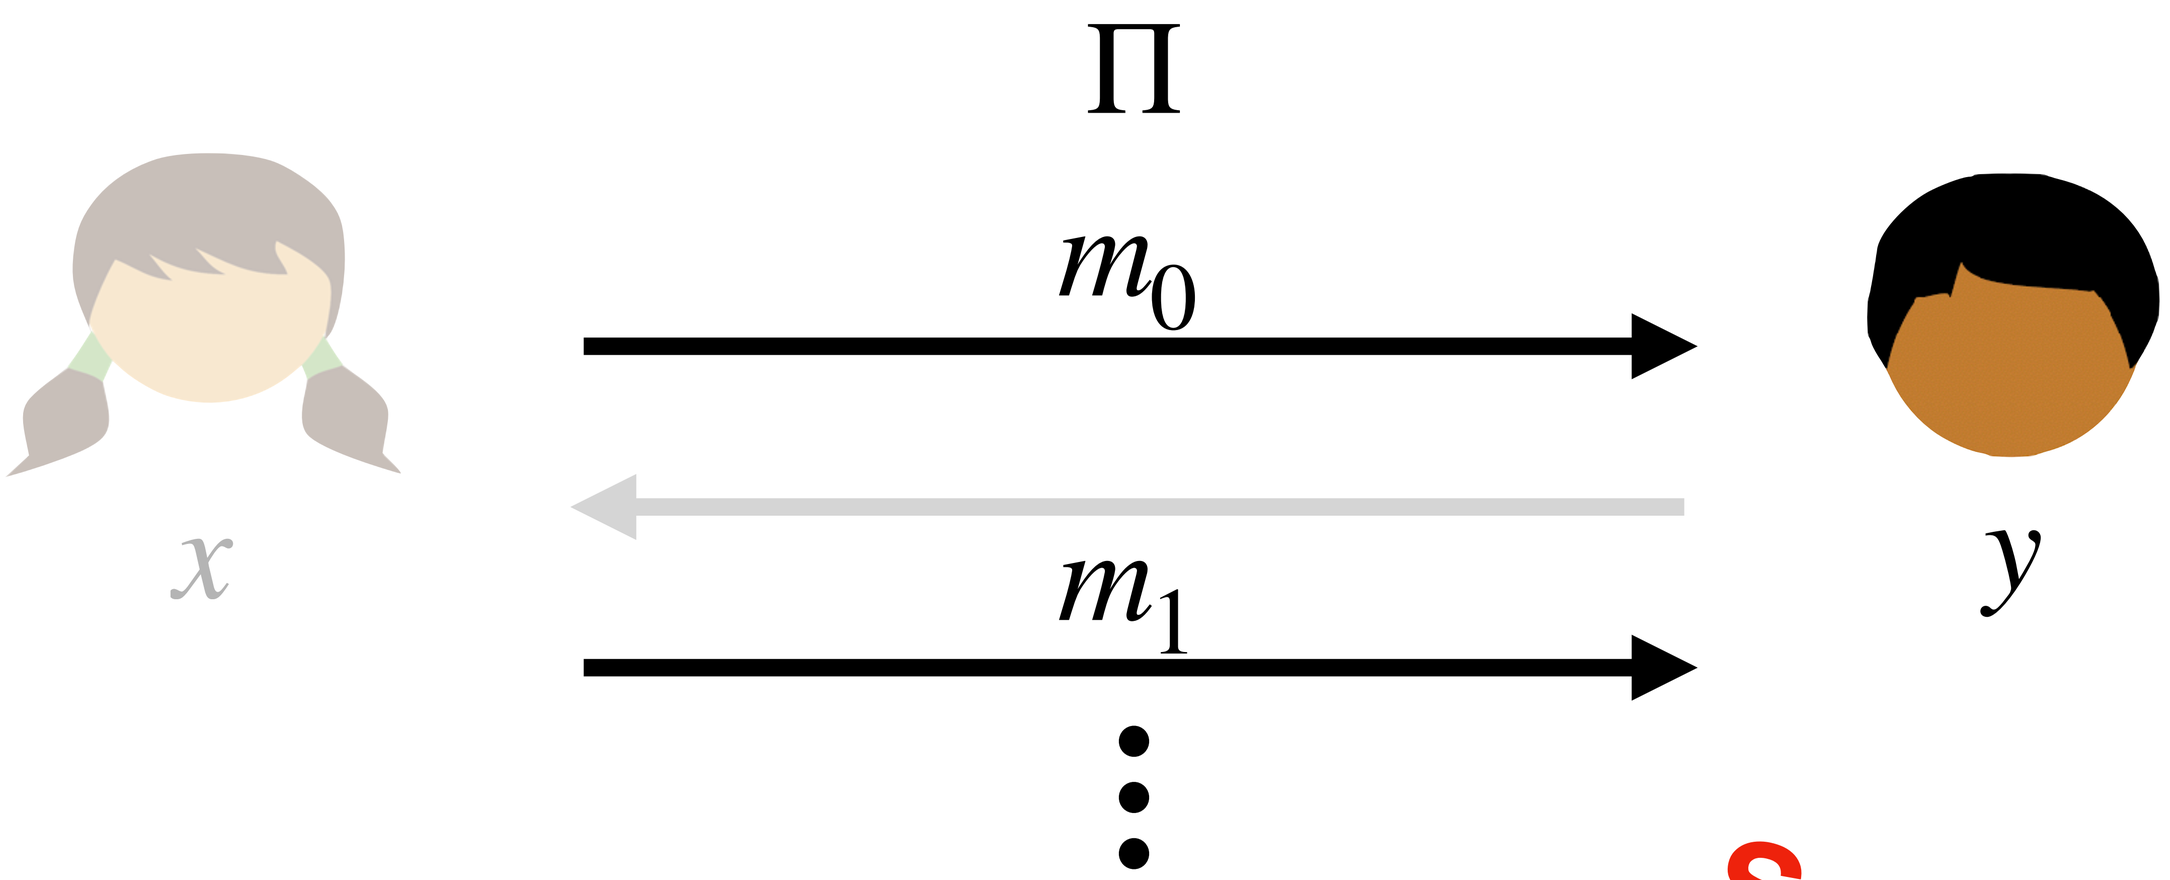

Security?

$$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$$

Real

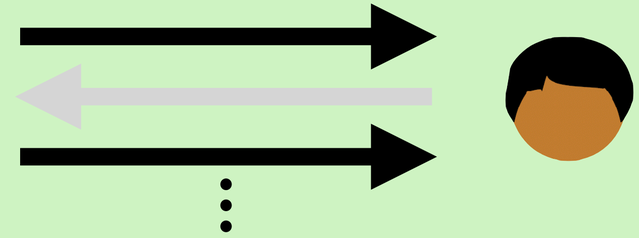

$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$

x

Trusted Third Party

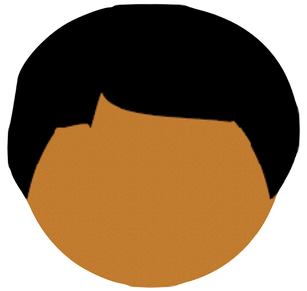

y

Real

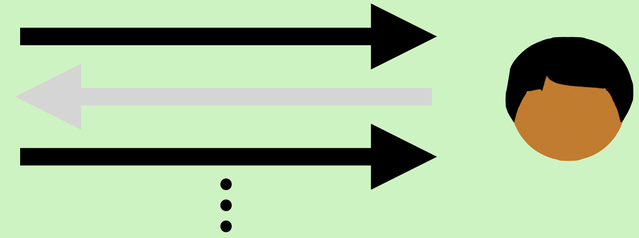

$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$

Real

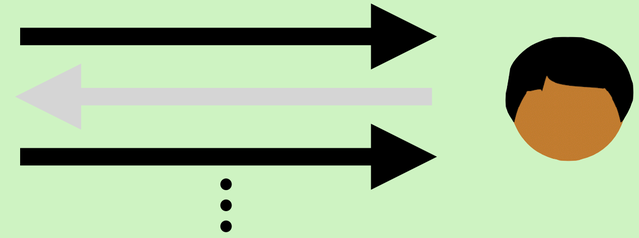

$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$

Real

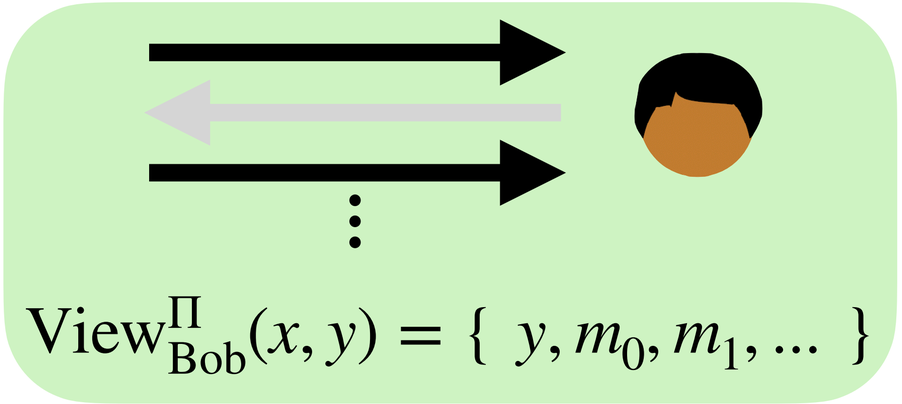

$$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$$

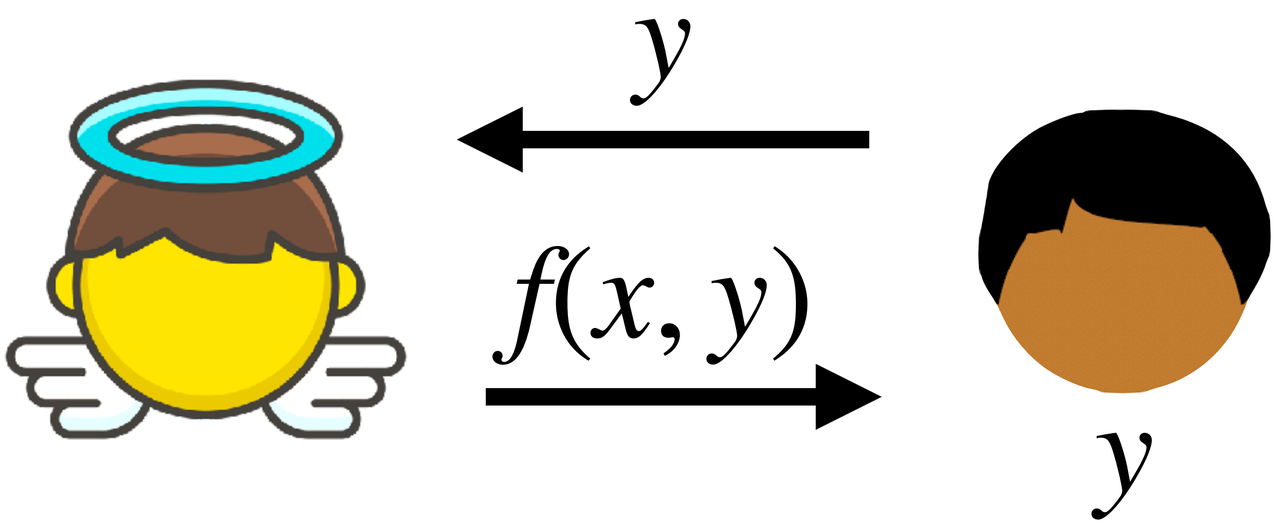

$$\text{View}^{\text{Ideal}}_{\text{Bob}}(x, y) = \{\ y, f(x, y)\ \}$$

Real

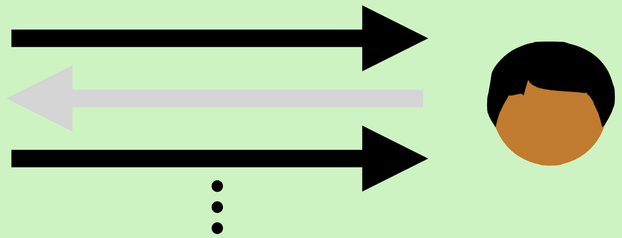

$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$

Ideal

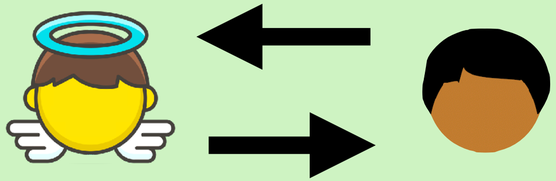

$\text{View}^{\text{Ideal}}_{\text{Bob}}(x, y) = \{\ y, f(x, y)\ \}$

These should “look the same”

Simulator

Real

Ideal

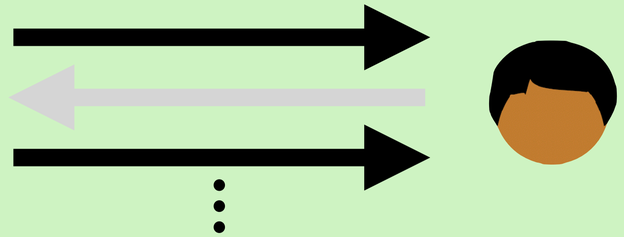

$\text{View}^{\Pi}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$

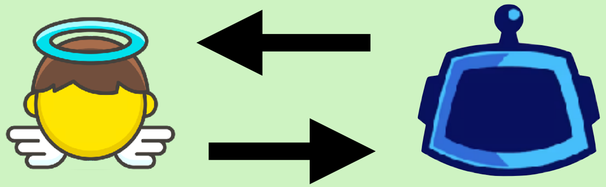

$\text{Output}^{\text{Sim}}_{\text{Bob}}(x, y) = \{\ y, m_0, m_1, \ldots\ \}$

These should “look the same”

Today's objectives

See applications of multiparty computation (MPC)

Sketch definition of ***semi-honest security***

Introduce the notion of a ***simulator***