

The Multiple Millionaires' Problem: New Algorithmic Approaches and Protocols

Tamir Tassa
The Open University
Ra'anana, Israel
tamirta@openu.ac.il

Avishay Yanai
Soda Labs
Tel Aviv, Israel
avishay@sodalabs.xyz

ABSTRACT

We study a fundamental problem in Multi-Party Computation, which we call the Multiple Millionaires' Problem (MMP). Given a set of private integer inputs, the problem is to identify the subset of inputs that equal the maximum (or minimum) of that set, without revealing any further information on the inputs beyond what is implied by the desired output. Such a problem is a natural extension of the Millionaires' Problem, which is the very first Multi-Party Computation problem that was presented in Andrew Yao's seminal work (FOCS 1982). A closely related problem is MaxP, in which the value of the maximum is sought. We study these fundamental problems and describe several algorithmic approaches and protocols for their solution. In addition, we compare the performance of the protocols under several selected settings. As applications of privacy-preserving computation are more and more commonly implemented in industrial systems, MMP and MaxP become important building blocks in privacy-preserving statistics, machine learning, auctions and other domains. One of the prominent advantages of the protocols that we present here is their simplicity. As they solve fundamental problems that are essential building blocks in various application scenarios, we believe that the presented solutions to those problems, and the comparison between them, will serve well future researchers and practitioners of secure distributed computing.

KEYWORDS

multi-party computation, millionaires' problem, the maximum problem, privacy-preserving computation

1 INTRODUCTION

Whenever the notion of Multi-Party Computation (MPC) is mentioned, it is almost always followed by a citation of Andrew Yao's seminal work from 1982, entitled "Protocols for secure computation" [31]. Yao presented there the first MPC protocol; the problem solved by that protocol was the Millionaires' Problem (MP): given two parties, P_1 and P_2 , each one holding a private nonnegative integer x_1 and x_2 , respectively, determine whether $x_1 < x_2$ or not, without revealing any further information on x_1 and x_2 .¹

¹This problem is sometimes referred to as the Greater Than problem.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.
Proceedings on Privacy Enhancing Technologies 2024(4), 784–796
© 2024 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2024-0141>



As MP is an essential building block for a vast array of MPC problems, many studies have presented over the years additional solutions to it, e.g. [3, 7, 9, 13, 14, 18, 20, 22, 26, 27]. We proceed to define a natural extension of MP:

DEFINITION 1. Let x_1, \dots, x_N , where $x_n \geq 0$ for $n \in [N]$, be private integers that are held by $K > 1$ distrustful parties. Let $M := \max\{x_n : n \in [N]\}$. Then the Multiple Millionaires' Problem (MMP) is the MPC problem of determining all indices $n \in [N]$ for which $x_n = M$.² The Maximum Problem (MaxP) is the MPC problem of computing $M = \max\{x_n : n \in [N]\}$.

In Definition 1 above we distinguish between the number of inputs, N , and the number of MPC parties, K . For example, if the MPC parties are the owners of the inputs, where party P_k holds n_k private inputs, then $N = \sum_{k=1}^K n_k$. In another possible scenario, the N inputs are held by external clients who distribute shares in them to K servers who act as the MPC parties.

As we show later, the two problems in Definition 1 are related: MMP can be reduced to MaxP with an additive overhead of linear size and constant depth (see Definitions 2-3), while MaxP can be reduced to MMP with an additive constant overhead of both size and depth.

The desired output in MMP is the *identities* of the "richest" members in some group of N parties ("millionaires"), and the goal is to carry out that computation in a manner that prevents those parties from learning any further information on the inputs of their peers, besides what they may infer from their own input and the computed output.

For the sake of convenience we focus here on the case of target functions that need to be maximized. However, our discussion also includes the case where the target function needs to be minimized. Namely, any algorithm for MMP can be easily used to find all indices $n \in [N]$ for which $x_n = m := \min\{x_n : n \in [N]\}$, and any algorithm for MaxP can be easily converted to an algorithm that computes the minimum m rather than the maximum M .

Identifying a set of indices that maximize (or minimize) some target function is a fundamental problem that is most useful in statistical analysis of data. In particular, it is an essential building block in machine learning, as many machine learning algorithms, such as k -means clustering [5, 24], k -nearest neighbors [8, 21], or fingerprint-matching [4, 19] include the identification of a subset of points in an Euclidean space that are *closest* to a given point in that space. (Note that the problem of finding the closest point is equivalent to MMP, by replacing each of the distances, say w , with $W - w$, where W is

²We note that MMP treats all N inputs similarly, while MP does not. Hence, when $N = 2$, MMP coincides with MP only when $x_1 \neq x_2$; if, however, $x_1 = x_2$, then MP will issue a negative answer that implies that $x_1 \geq x_2$, while MMP will issue the accurate answer that $x_1 = x_2$.

some upper bound on all distances.) For example, in k -means clustering one is given a set of k -means $\mathcal{M} := \{m_1, \dots, m_k\} \subset \mathbb{R}^d$, and a set of observations $\mathcal{X} := \{x_1, \dots, x_N\} \subset \mathbb{R}^d$, and for each given observation $x_n \in \mathcal{X}$, it is needed to find the mean (or means) in \mathcal{M} that minimize the distance to x_n . In *privacy-preserving* machine learning, and in particular in federated (collaborative) learning, the observations need to remain secret; in such a case, all computations must be carried out in a privacy-preserving manner and, therefore, one has to solve an instance of MMP. In the k -nearest neighbors algorithm, on the other hand, the goal is to find the k points out of a set of N points, $\mathcal{X} = \{x_1, \dots, x_N\} \subset \mathbb{R}^d$, that are closest to a new observation $y \in \mathbb{R}^d$; namely, the sought-after points here are the k points that issue the smallest values for the target (distance) function, $f(x) := \|x - y\|$. When the k -nearest neighbors algorithm has to be applied in a privacy-preserving manner, the computational problem that it has to solve may be reduced to a sequence of instances of MMP. Specifically, we run MMP once, and identify the points in \mathcal{X} that are closest to y . Assume that this execution of MMP identified $\ell \geq 1$ such points. If $\ell \geq k$ we are done. Otherwise, We remove those ℓ points from \mathcal{X} and repeat the process until we have found k points in \mathcal{X} that are closest to y .

A different application where MMP plays a crucial role is that of privacy-preserving (sealed-bid) auctions. There too, it is needed to identify a set of parties (bidders) that maximize a target function (the bid), without revealing any further information on the inputs (bids). See the seminal work of Naor et al. [25] and the recent survey [2] and the references therein.

The cost of privacy-preserving computation can be roughly separated into two metrics: communication and latency (being the time interval from submitting the inputs until obtaining the output). Typically, there is a tradeoff between the two, as expressed by the solutions presented in this paper (see Section 7.3). Specifically, we study the tradeoff between the size of a protocol and its depth (see Definitions 2-3). The communication-latency tradeoff is relevant in particular to GMW-style MPC protocols [16], in which the parties interact for every non-linear gate, and to leveled fully homomorphic encryption (FHE) based protocols, in which circuit evaluation is done over ciphertexts. In the former, the latency is affected mainly by the depth of the protocol, as the parties interact in every non-linear operation. In the latter, the depth has an impact on the length of the keys with which the inputs are encrypted, as well as on the computation time for every operation. Indeed, deep circuits imply larger encryption keys, what implies larger ciphertexts, which are presented as polynomials of larger degree, and, consequently, multiplying polynomials incurs larger runtimes.

1.1 Our contributions and outline

Even though MMP is a very natural extension of the (two-party) Millionaires' Problem, very few studies so far devised MPC protocols for solving it or other related problems (see the survey of related work in Section 8.) In this study we propose several algorithmic approaches towards its solution as well as concrete protocols.

We begin with a discussion of essential preliminaries in Section 2: in Section 2.1 we present a commonly used abstraction method that allows us to present and analyze secure computation protocols in a general context that is indifferent to the underlying setting

and the implementation details; then, in Section 2.2, we present the complexity measures for assessing the performance of such protocols. In Section 3 we present a classical binary-tree-based protocol (Protocol 1) and a generalization to any degree (Protocol 4); these protocols are based on a "tournament" between the "candidates" (inputs) to determine the "winners" (maximal inputs), through several rounds of comparisons. To the best of our knowledge, the general-degree tree-based protocol (Protocol 4) is new, whereas the binary-tree based-protocol was already proposed and used in the literature, e.g. [12, 24]. In Section 4 we present novel constant-depth protocols for solving MaxP and MMP (Protocols 5 and 6) that are tailored to cases where the inputs are drawn from a small domain. Then, we present in Section 5 protocols that are based on the inputs' decomposition into bits (Protocol 7) or into digits in larger bases (Protocol 8). In Section 6 we present Protocol 11 for solving MaxP and MMP. It is inspired by the protocols of [1] for finding the k -th ranked element in a union of private datasets (namely, the value in the k -th entry in the unified dataset when it is sorted in an increasing order). We evaluate the performance of all of those protocols in Section 7; the analysis is based on an implementation of passively secure MPC protocols, which allows ignoring 'noisy' costs, like consistency checks of inputs, that are typically added to actively secure protocols. Furthermore, the asymptotic analysis of the cost of passively secure protocols captures the actively secure ones as well, as there are efficient compilers from passive to active security (e.g., [15]). We review related work in Section 8 and conclude in Section 9.

A prominent advantage of the presented protocols is their simplicity. Since MaxP and MMP are essential building blocks in important applications, the introduction of novel algorithmic approaches and protocols for those problems, and the comparison between the presented solutions, will serve well future researchers of MPC and practitioners of secure distributed computing.

2 PRELIMINARIES

2.1 The Arithmetic Black-Box (ABB)

Practical protocols for secure computation use a circuit representation of the functionality that they compute. The gates in the circuit may be arithmetic, where computation is carried out by addition and multiplication gates over an agreed upon field, or Boolean, where computation is carried out by AND and XOR gates. Note that Boolean gates adhere to arithmetics over \mathbb{Z}_2 , thus, they can use the arithmetic notation ('+', '.') as well.

Typically, privacy-preserving protocols abstract out the underlying techniques through an abstraction called an "Arithmetic Black Box" (ABB). An ABB is an ideal functionality that can be realized in many ways, depending on the setting (e.g., network topology and synchrony assumptions) and the security needs (i.e., semi-honest/malicious, static/adaptive adversaries, with/without abort, perfect/statistical/computational security, and different collusion scenarios). The privacy-preserving protocol directly inherits these properties from the sub-protocol that realizes the ABB. This way, it is possible to describe a privacy-preserving protocol that can fit several settings at the same time, given a suitable realization of the ABB functionality in the protocol. Furthermore, by relying

on the ABB functionality, one does not have to analyse the communication and computation complexity of its operations, as they depend on the specific realization. The ABB ideal functionality is given in Functionality 1. Abstraction using an ideal functionality is a common practice in the art of secure protocol design, which enables focusing on the logic of the application rather than on the underlying infrastructure.

In the beginning of the protocol the parties initialize the ABB functionality with the field \mathbb{F} over which all computations take place. We assume hereinafter that $\mathbb{F} = \mathbb{Z}_p$ for a sufficiently large prime p that is greater than all inputs of the MMP problem. After that initialization the parties may use any interface listed in Functionality 1 as many times as needed, as part of the privacy-preserving protocol. See [6] for the syntax and usability of ideal functionalities and [10] for the introduction of the arithmetic black-box.

In Functionality 1 we use the term *handle*. A handle can be seen as a public pointer to the secret held by the ABB, so that the parties can relate to it in their calls. A possible instantiation of a handle is as follows. On inputting a secret or a public value x , the ABB picks a unique identifier (e.g., a long random string) and outputs it as x 's handle. In addition, when the parties invoke a computation on the secrets represented by handles h_1, \dots, h_k the ABB performs the computation on the secrets and picks a new unique identifier for the result.

In a protocol description, instead of explicitly mentioning that the parties send some message through the ABB interface, it is convenient to use a shorthand. Specifically:

- instead of writing that party P_n calls (PrivateInput, n, a) and all other parties call (PrivateInput, n) so that all parties obtain $[a]$, we write simply $[a] \leftarrow (\text{PrivateInput}, n, a)$.
- instead of writing that all parties call (Add, $[a], [b]$) and obtain the handle $[c]$, we write $[c] \leftarrow [a] + [b]$;
- a call to (Multiply, $[a], [b]$) is written $[c] \leftarrow [a] \cdot [b]$;
- a call to (Duplicate, $[a]$) which returns to all parties a handle $[b]$ with the same value as $[a]$ is written $[b] \leftarrow [a]$;
- a call to (AffineComb, $[a_1], \dots, [a_k], c_0, c_1, \dots, c_k$) is written $[c] \leftarrow c_0 + c_1 \cdot [a_1] + \dots + c_k \cdot [a_k]$;
- a call to (Inverse, $[a]$) is written $[c] \leftarrow [a^{-1}]$;
- a call to (Compare, $[a], [b]$) is replaced with the handle that it returns, i.e. $[1_{a < b}]$.³
- a call to (Equal, $[a], [b]$) is replaced with the result $1_{a=b}$.
- a call to (OR $_m$, $[a_1], \dots, [a_m]$) is written $[c] \leftarrow \bigvee_{i=1}^m [a_i]$.

There are plenty of potential instantiations for the ABB functionality, depending on the number of parties (e.g., it is common to have more efficient protocols for a small number of parties, say $N \in \{2, 3, 4, 5\}$), the adversarial model (e.g., protocols that protect against a semi-honest and static adversary who can corrupt only a few parties are obviously simpler and more efficient than those protecting against an active and adaptive adversary who may corrupt almost all parties), the network topology (e.g., it is easier to communicate over a full mesh rather than over a sparse network), and the underlying setting.

We conclude by noting that some of the interfaces in the ABB functionality could be realized by other ones (see more on that in [30, Appendix A.1]).

³If \mathcal{P} is a predicate then $1_{\mathcal{P}}$ is a bit that equals 1 if the predicate holds and 0 otherwise.

FUNCTIONALITY 1. (*Arithmetic Black-Box*)

Initialize. On input (Init, \mathbb{F}) from all parties, store the field \mathbb{F} .

Public Input. [PuI] On input (PublicInput, a) from all parties, where $a \in \mathbb{F}$, store a and return a handle $[a]$ to the parties.

Private Input. [PrI] On input (PrivateInput, n, a) from party P_n , where $a \in \mathbb{F}$, and input (PrivateInput, n) from all other parties, store a and return a handle $[a]$ to the parties.

Public Output. [PuO] On input (PublicOutput, $[a]$) from all parties, reveal a to all parties.

Private Output. [PrO] On input (PrivateOutput, $n, [a]$) from all parties, reveal a to party P_n .

Random. [RND] On input (Random) from all parties, sample $r \leftarrow \mathbb{F}$ uniformly and return the handle $[r]$ to the parties.

Addition. On input (Add, $[a], [b]$) from all parties, compute $c = a + b$ in \mathbb{F} and return a handle $[c]$ to all parties.

Multiplication. [MUL] On input (Multiply, $[a], [b]$) from all parties, compute $c = a \cdot b$ in \mathbb{F} and return a handle $[c]$ to all parties.

Duplicate. On input (Duplicate, $[a]$) from all parties, the functionality generates a new handle $[b]$ to a secret b where $b = a$, and returns the handle $[b]$ to all parties.

Affine combination. On input (AffineComb, $[a_1], \dots, [a_k], c_0, c_1, \dots, c_k$) from all parties, compute $c = c_0 + \sum_{i=1}^k c_i a_i$ in \mathbb{F} and return a handle $[c]$ to all parties.

Inverse. [INV] On input (Inverse, $[a]$) from all parties, where $a \in \mathbb{F} \setminus \{0\}$, compute $c = a^{-1}$ and return a handle $[c]$ to all parties.

Equality. [EQV] On input (Equal, $[a], [b]$) from all parties, compute and return $c \leftarrow 1_{a=b}$.

Comparison. [COM] On input (Compare, $[a], [b]$) from all parties, compute $c \leftarrow 1_{a < b}$ and return a handle $[c]$ to all parties. (We focus here on prime-ordered fields $\mathbb{F} = \mathbb{Z}_p$ in which $a < b$ means that a is smaller than b when both a and b are viewed as integers.)

OR. [OR] On input (OR $_m$, $[a_1], \dots, [a_m]$) from all parties, if $a_i \in \{0, 1\}$ for all $i \in [m]$, compute $c \leftarrow \bigvee_{i=1}^m a_i$ and return a handle $[c]$ to all parties; otherwise, return \perp to all parties.

2.2 Measures of protocols' complexity

Hereinafter, when analyzing the complexity of a given MPC protocol we will refer to the protocol's size and depth, where the *size* addresses the protocol's communication complexity and the *depth* addresses its round complexity. Each MPC protocol induces a tree of calls to the ABB functionality. Let \mathcal{I} be the set of all interfaces in the ABB Functionality 1. Then each node in the tree that is induced by a given protocol is labeled by some $A \in \mathcal{I}$. All nodes in the same distance from the root node represent calls that can be executed in parallel. On the other hand, a pair of a father node with a son node represents a pair of calls to the ABB functionality where the latter call depends on the completion of the former.

DEFINITION 2 (PROTOCOL'S SIZE). *Given an MPC protocol Π we let c_A be the number of invocations of the interface $A \in \mathcal{I}$ during Π 's execution. Then the size of protocol Π is defined as $\text{size}(\Pi) = \sum_{A \in \mathcal{I}} c_A \cdot \text{size}(A)$, where $\text{size}(A)$ is the message complexity of the invocation of A in a specific implementation of the ABB functionality.*

DEFINITION 3 (PROTOCOL'S DEPTH). Let $w := (A_1, \dots, A_d)$ be a path in the tree of calls that protocol Π induces. Its depth is defined as $\text{depth}_w := \sum_{i=1}^d \text{depth}(A_i)$, where $\text{depth}(A_i)$ is the round complexity of the invocation of A_i in a specific implementation of the ABB functionality. Then the depth of protocol Π is defined as $\max_w \text{depth}_w$ where the maximum is over all paths in the tree that Π induces.

In our analysis, we break the depth of the implementation of each interface A to its part that is independent of the input and its part that does depend on the input. For an interface A , we call the former the *preprocessing depth* and denote it by $\text{preDepth}(A)$, while the latter is called the *online depth* and is denoted $\text{onDepth}(A)$. Similarly, the depth of a protocol Π is also split into a preprocessing depth and an online depth, $\text{depth}(\Pi) = \text{preDepth}(\Pi) + \text{onDepth}(\Pi)$. Note that if protocol Π invokes a set of interfaces $\mathcal{J} \subseteq \mathcal{I}$, then $\text{preDepth}(\Pi) = \max\{\text{preDepth}(A) : A \in \mathcal{J}\}$ because the preprocessing part for all interface invocations can be done simultaneously. As for the online depth, the exact depth depends on the invocation pattern in each protocol.

When analyzing the size and depth of our protocols, we use abbreviated notations as listed in Functionality 1. So, for example, the notation $[\text{COM}]$ denotes the *size* of the Compare interface, when speaking about the size of a protocol, while it denotes the *preprocessing depth* (resp., *online depth*) of the interface when discussing the protocol's preprocessing depth (resp., online depth). In such analysis we ignore calls to Add, AffineComb and Duplicate, since they incur no communication at all.

3 TREE-BASED PROTOCOLS

In this section we present protocols that solve MMP by decomposing the original problem, that involves many inputs, into a hierarchy (tree) of smaller instances of MMP, each one consisting of a smaller number of inputs. In Section 3.1 we describe a protocol due to [12, 24] that is based on a binary tree; in that protocol, each smaller instance of MMP involves just two inputs. Then, in Section 3.2, we present a generalized protocol in which the smaller MaxP instances involve up to N' (where $2 \leq N' \leq N$) inputs.

3.1 A binary tree-based protocol

Here we present a basic MMP protocol that reduces the Multiple Millionaires' Problem into the (Two) Millionaires' Problem MP. Protocol 1 is executed by the parties P_1, \dots, P_N towards identifying all maximal inputs. The protocol uses a binary tree based comparison in order to compute the maximal value.

Initially, all parties input their secret inputs, x_1, \dots, x_N (Lines 1-2). The computation that follows is performed in L rounds, where $L = \lceil \log N \rceil$ (hereinafter, unless otherwise stated, $\log = \log_2$). Denote by N_ℓ the number of candidates in the ℓ -th round, $\ell \in [L]$, and let $x_1^\ell, \dots, x_{N_\ell}^\ell$ be those candidates. In the first round the candidates are $x_n^1 = x_n$, $n \in [N]$, so N_1 is initialized to N (Line 3).

As mentioned, the computation of the maximal value M is done in L rounds (Lines 4-10). In the ℓ -th round every two consecutive candidates, x_k^ℓ and x_{k+1}^ℓ , are compared (Line 6). The comparison result is used in order to compute the maximum of those two candidates using the equality $\max\{a, b\} = a + 1_{a < b} \cdot (b - a)$; it is then

stored in $[x_{(k+1)/2}^{\ell+1}]$ and percolated to the next round (Line 7). If the number of candidates in the ℓ -th round, N_ℓ , is odd, then the last candidate ($x_{N_\ell}^\ell$) is directly percolated to the next round (Lines 8-9). Finally, the number of values in the next layer is updated (Line 10). After completing the L rounds, the maximum M is the single value in the $(L + 1)$ -th layer (Line 11). The protocol ends with a testing the equality of each of the inputs to the computed maximum (Lines 12-13) and outputting the indices of all maximal inputs.

Protocol 1: A binary-tree-based protocol for MMP

Parameters: N - number of inputs; $L = \lceil \log N \rceil$.

Private Inputs: $x_1, \dots, x_N \in \mathbb{F}$.

```

1 forall  $n \in [N]$  do
2    $[x_n^1] \leftarrow (\text{PrivateInput}, n, x_n)$ 
3  $N_1 \leftarrow N$ .
4 forall  $\ell = 1, 2, \dots, L$  do
5   forall Odd  $k$  s.t.  $k < N_\ell$  do
6     The parties compute  $[1_{x_k^\ell < x_{k+1}^\ell}]$ .
7      $[x_{(k+1)/2}^{\ell+1}] \leftarrow [x_k^\ell] + [1_{x_k^\ell < x_{k+1}^\ell}] \cdot ([x_{k+1}^\ell] - [x_k^\ell])$ .
8   if  $N_\ell$  is odd then
9      $[x_{N_\ell/2+1}^{\ell+1}] \leftarrow [x_{N_\ell}^\ell]$ .
10   $N_{\ell+1} \leftarrow \lceil N_\ell/2 \rceil$ .
11  $[M] \leftarrow [x_1^{L+1}]$ .
12 forall  $n \in [N]$  do
13   The parties compute  $1_{x_n^1 = M}$ .
Output: The indices  $n \in [N]$  for which
 $x_n = M := \max\{x_i : i \in [N]\}$ .

```

The protocol consists of N calls to PrivateInput (Line 2), followed by $N - 1$ calls to Compare and Multiply (Line 7). Afterwards, the protocol performs N calls to Equality (Line 13) and PublicOutput. Note that all calls to PrivateInput occur in parallel, as well as the calls to Equality and PublicOutput. In addition, all calls to Compare and Multiply occur in $\lceil \log N \rceil$ iterations. Therefore, the size and depth of Protocol 1 are:

$$\begin{aligned} \text{size}(\text{Protocol 1}) &= N \cdot [\text{PR}] + (N - 1) \cdot [\text{COM}] + \\ &\quad (N - 1) \cdot [\text{MUL}] + N \cdot [\text{EQU}] + N \cdot [\text{PUO}], \\ \text{preDepth}(\text{Protocol 1}) &= \max([\text{PR}], [\text{COM}], [\text{MUL}], [\text{EQU}], [\text{PUO}]), \\ \text{onDepth}(\text{Protocol 1}) &= [\text{PR}] + \lceil \log N \rceil \cdot ([\text{COM}] + [\text{MUL}]) + \\ &\quad [\text{EQU}] + [\text{PUO}]. \end{aligned}$$

3.2 A protocol based on higher degree trees

Here we present Protocol 4 for solving MMP. That protocol is also based on a comparison tree, only that here the degrees of all nodes are (at most) some parameter $N' \geq 2$. The depth of Protocol 4 is $\lceil \log_{N'} N \rceil$. If we take $N' = 2$ we restore Protocol 1. If, on the other hand, we take $N' = \Theta(\sqrt{N})$, we get a shallow tree with a depth of $O(1)$.

Before presenting Protocol 4, we present two sub-protocols that it invokes, for solving MMP and MaxP in constant depth. Sub-protocol 2 securely computes the indices of the maximal value. It receives as input handles to N' candidate values. First, it computes in parallel the comparison bits between every pair of inputs (Line 1-2). It

then uses those bits to compute a bit for each input that indicates whether that input is maximal (Lines 3-4).

Sub-protocol 2: An all-to-all-comparison protocol for MMP

Inputs: Handles $[x_1], \dots, [x_{N'}]$ where $x_i \in \mathbb{F}$.

1 **forall** $1 \leq i \neq j \leq N'$ **do**

2 | $[b_{i,j}] \leftarrow [1_{x_i < x_j}]$.

3 **forall** $i \in [N']$ **do**

4 | $[b_i] \leftarrow 1 - \bigvee_{j \in [N], j \neq i} [b_{i,j}]$.

Output: The handles $([b_1], \dots, [b_{N'}])$.

Sub-protocol 3 computes the maximal value. It too receives as input handles to N' candidate values. It then invokes Sub-protocol 2 and gets handles to the bits $[b_i]$, $i \in [N']$, that indicate the maximality of each input (Line 1). It then proceeds to add up all maximal inputs into \hat{M} and count them into \hat{N} (Lines 2-3). The maximum is obtained by multiplying \hat{M} with the inverse of \hat{N} .

Sub-protocol 3: A sub-protocol for MaxP

Inputs: Handles $[x_1], \dots, [x_{N'}]$ where $x_i \in \mathbb{F}$.

1 $([b_1], \dots, [b_{N'}]) \leftarrow \text{MMP}([x_1], \dots, [x_{N'}])$.

2 $[\hat{M}] \leftarrow \sum_{i \in [N']} ([b_i] \cdot [x_i])$.

3 $[\hat{N}] \leftarrow \sum_{i \in [N']} [b_i]$.

4 $[M] \leftarrow [\hat{N}^{-1}] \cdot [\hat{M}]$.

Output: The handle $[M]$ to $M = \max\{x_i : i \in [N']\}$.

We now turn to Protocol 4. Lines 1-3 in it are the same as in Protocol 1. Lines 4-7 are equivalent to Lines 4-10 in Protocol 1, with the only difference being the fact that in the ℓ -th iteration, $1 \leq \ell \leq \lceil \log_{N'} N \rceil$, the N_ℓ candidates are grouped into groups of at most N' candidates in each (instead of 2 in Protocol 1) and then the maximum of each such group is percolated to the next iteration. The conclusion of the protocol (Lines 8-10) is similar to the conclusion of Protocol 1 (Lines 11-13 there).

Note that if we wish to solve MaxP rather than MMP, we can run Lines 1-8 in Protocol 4 and then output M .

We conclude with analyzing the costs of those protocols. Sub-protocol 2's size is $((N')^2 - N') \cdot [\text{COM}] + N' \cdot [\text{OR}]_{N'}$, where the subscript N' in $[\text{OR}]_{N'}$ denotes the number of operands in the OR operation. Sub-protocol 2's preprocessing and online depths are $\max([\text{COM}], [\text{OR}]_{N'})$ and $[\text{COM}] + [\text{OR}]_{N'}$, respectively. These costs of Sub-protocol 2 imply that the costs of Sub-protocol 3 are:

$$\text{size}(\text{Sub-protocol 3}) = ((N')^2 - N') \cdot [\text{COM}] + N' \cdot [\text{OR}]_{N'} + (N' + 1) \cdot [\text{MUL}] + [\text{INV}],$$

$$\text{preDepth}(\text{Sub-protocol 3}) = \max([\text{COM}], [\text{OR}]_{N'}, [\text{MUL}], [\text{INV}]),$$

$$\text{onDepth}(\text{Sub-protocol 3}) = [\text{COM}] + [\text{OR}]_{N'} + 2[\text{MUL}] + [\text{INV}].$$

Letting $[\text{MAXP}]$ denote the costs of Sub-protocol 3, we have

$$\text{size}(\text{Protocol 4}) = N \cdot [\text{PRI}] + T(N, N') \cdot [\text{MAXP}] + N \cdot [\text{EQU}],$$

where $T(N, N') = \sum_{\ell=1}^{L'} [N_\ell / N']$ and N_ℓ is defined recursively in Lines 3 & 7. (Note that if $N = (N')^{L'}$ the tree is regular, in the

Protocol 4: A general-tree-based protocol for MMP

Parameters: N - number of inputs; N' , where $2 \leq N' \leq N$
 - number of inputs that are compared in one call to Sub-protocol 3 (MaxP); $L' = \lceil \log_{N'} N \rceil$.

Private Inputs: $x_1, \dots, x_N \in \mathbb{F}$.

1 **forall** $n \in [N]$ **do**

2 | $[x_n^1] \leftarrow (\text{PrivateInput}, n, x_n)$

3 $N_1 \leftarrow N$.

4 **forall** $\ell = 1, 2, \dots, L'$ **do**

5 | **forall** k s.t. $k = 1 \pmod{N'}$ and $k \leq N_\ell$ **do**

6 | | $[x_{1+(k-1)/N'}^{\ell+1}] \leftarrow \text{MaxP}([x_k^\ell], [x_{k+1}^\ell], \dots, x_{\min\{k+N'-1, N_\ell\}}^\ell)$.

7 | $N_{\ell+1} \leftarrow \lceil N_\ell / N' \rceil$.

8 $[M] \leftarrow [x_1^{L'+1}]$.

9 **forall** $n \in [N]$ **do**

10 | The parties compute $1_{x_n^1=M}$. (Recall that x_n^1 is the n -th input).

Output: The indices $n \in [N]$ for which

$$x_n = M := \max\{x_i : i \in [N]\}.$$

sense that all non-leaf nodes have an out-degree N' , and then $T(N, N') = \frac{N-1}{N'-1}$. Protocol 4's depths are:

$$\begin{aligned} \text{preDepth}(\text{Protocol 4}) &= \max([\text{PRI}], [\text{MAXP}], [\text{EQU}]), \\ \text{onDepth}(\text{Protocol 4}) &= [\text{PRI}] + L' \cdot [\text{MAXP}] + [\text{EQU}]. \end{aligned}$$

4 CONSTANT-DEPTH PROTOCOLS FOR SMALL DOMAINS

We begin this section by presenting a constant-depth protocol for solving MaxP (Section 4.1). Then we use that protocol as a basis for a constant-depth protocol for solving MMP (Section 4.2). The size of the protocols is proportional to the underlying field size and, therefore, they are relevant in cases where there is a small known upper bound on the inputs.

4.1 A protocol for solving MaxP

Protocol 5 starts with each P_n , $n \in [N]$, submitting its input, which is an integer $x_n \in [0, Q]$, for some publicly known upper bound Q , in the form of $Q-1$ bits $y_{n,i}$ that represent it (Lines 1-3). Specifically, each input $x \in [0, Q]$ is represented by a sequence of $Q-1$ bits in which the first x bits are 1 and the proceeding ones are all 0; it is that form of input encoding that restricts the scope of this protocol to small domains. Then, the parties compute a similar representation of M into the vector $\mathbf{z} = (z_1, \dots, z_{Q-1})$, using the \bigvee operator (Lines 4-5). The \bigvee operator can be computed by a constant-depth sub-protocol (see [30, Appendix A.2.3]). Finally, the maximum M is computed by adding all entries in \mathbf{z} (Line 6).

The protocol performs $N(Q-1)$ parallel invocations of PrivateInput (Lines 1-3) and $Q-1$ parallel invocations of Or (Lines 4-5). Hence, its costs are:

$$\begin{aligned} \text{size}(\text{Protocol 5}) &= N(Q-1) \cdot [\text{PRI}] + (Q-1) \cdot [\text{OR}]_N, \\ \text{preDepth}(\text{Protocol 5}) &= \max([\text{PRI}], [\text{OR}]_N), \\ \text{onDepth}(\text{Protocol 5}) &= [\text{PRI}] + [\text{OR}]_N. \end{aligned}$$

Protocol 5: A monotone-representation-based protocol for MaxP

Parameters: N - number of inputs; Q - an upper bound on the inputs (each input is in $[0, Q - 1]$).

Private Inputs: $x_1, \dots, x_N \in [0, Q]$.

```

1 forall  $n \in [N]$  do
2   | forall  $i = 1, \dots, Q - 1$  do
3     |    $[y_{n,i}] \leftarrow (\text{PrivateInput}, n, y_{n,i})$ , where  $y_{n,i} = 1_{x_n \geq i}$ 
4   forall  $i = 1, \dots, Q - 1$  do
5     |    $[z_i] \leftarrow \bigvee_{n \in [N]} [y_{n,i}]$ 
6    $[M] \leftarrow \sum_{i=1}^{Q-1} [z_i]$ 
Output: A handle  $[M]$  to  $M := \max\{x_n : n \in [N]\}$ .

```

4.2 A protocol for solving MMP

Here we present Protocol 6 – an MMP protocol that is based on the MaxP Protocol 5. It starts with an invocation of Protocol 5 to compute the maximum M of all inputs (Line 1). Then, the parties go over all inputs and test their equality to M (Lines 2-3). The protocol's output identifies all inputs that equal M .

Protocol 6: A monotone-representation-based protocol for MMP

Parameters: N, Q - two positive integers, denoting the number of inputs and an upper bound on them.

Private Inputs: $x_1, \dots, x_N \in [0, Q]$;

```

1 The parties compute  $[M] \leftarrow [\text{MaxP}\{x_1, \dots, x_N\}]$ .
2 forall  $n \in [N]$  do
3   | The parties compute  $1_{x_n=M}$ .

```

Output: The indices $n \in [N]$ for which $x_n = M := \max\{x_i : i \in [N]\}$.

Protocol 6 consists of a call to Protocol 5 to compute M , followed by N parallel procedures of testing the equality of each input to M (Lines 2-5). Hence, its costs are as follows:

$$\begin{aligned} \text{size}(\text{Protocol 6}) &= \text{size}(\text{Protocol 5}) + N \cdot [\text{EQV}] + N \cdot [\text{PUO}], \\ \text{preDepth}(\text{Protocol 6}) &= \max(\text{preDepth}(\text{Protocol 5}), [\text{EQV}], [\text{PUO}]), \\ \text{onDepth}(\text{Protocol 6}) &= \text{onDepth}(\text{Protocol 5}) + [\text{EQV}] + [\text{PUO}]. \end{aligned}$$

By plugging the size and the preprocessing and online depths of Protocol 5 as derived in Section 4.1, we arrive at the final size and depths of Protocol 6:

$$\begin{aligned} \text{size}(\text{Protocol 6}) &= N(Q - 1) \cdot [\text{PRI}] + (Q - 1) \cdot [\text{OR}]_N + \\ &\quad N \cdot [\text{EQV}] + N \cdot [\text{PUO}]; \\ \text{preDepth}(\text{Protocol 6}) &= \max([\text{PRI}], [\text{OR}]_N, [\text{EQV}], [\text{PUO}]), \\ \text{onDepth}(\text{Protocol 6}) &= [\text{PRI}] + [\text{OR}]_N + [\text{EQV}] + [\text{PUO}]. \end{aligned}$$

5 PROTOCOLS BASED ON THE INPUTS' DIGIT DECOMPOSITION

Here we present two protocols for solving MMP that are based on the digits of the inputs' representation in some number base. We begin with Protocol 7 (Section 5.1) that uses a binary representation of the inputs. We then present Protocol 8 (Section 5.2) that reduces

depth by considering a representation of the inputs in a 2^d -base, for some $d > 1$.

5.1 Binary representation of inputs

Protocol 7 gets as inputs N integers, x_1, \dots, x_N . As stated earlier, we assume that all inputs are smaller than some prime p , where the underlying field is $\mathbb{F} = \mathbb{Z}_p$. Hereinafter we let $B = \lceil \log p \rceil$ denote the number of bits in the binary representation of elements in \mathbb{F} . Hence, the input x_n , $n \in [N]$, has a binary representation by B bits, $(x_{n,B-1}, \dots, x_{n,0})$, meaning that $x_{n,b} \in \{0, 1\}$ and $x_n = \sum_{b=0}^{B-1} x_{n,b} 2^b$. In addition, we define $x_{n,B} = 1$ for all $n \in [N]$. The protocol outputs an N -dimensional Boolean vector that identifies the maximal inputs.

The protocol starts by submitting all input bits (Lines 1-3). It then iterates over the input bits from bit number $b = B - 1$, the MSB, to bit number $b = 0$, the LSB (Lines 4-7). We refer to the collection of all bits in the same bit position, $\{x_{n,b} : n \in [N]\}$, as the b -th bit column. For each of those bit columns the protocol computes $s_b := \bigvee_{n \in [N]} (x_{n,b+1} \cdot x_{n,b})$ (Line 5). If $s_b = 1$ then there is at least one input x_n for which $x_{n,b} = 1$ and, in addition, the current value of $x_{n,b+1}$ (after it may have been updated in the previous iteration) also equals 1; in that case, the protocol updates each entry in that column to equal the product of that entry and the preceding one. Otherwise, if $s_b = 0$, the protocol updates the b -th bit column to be the same as the preceding bit column. This computation is done in Lines 6-7. (Note that all computations are done in \mathbb{Z}_2 so that the minus and plus operations are just XORs). At the conclusion of this loop, the protocol outputs the bits in the 0-th bit column (Line 8) since those bits identify all maximal inputs, as we claim next:

LEMMA 4. $x_{n,0} = 1$ if and only if $x_n = M = \max\{x_i : i \in [N]\}$.

Proof. For each $b = B - 1, \dots, 0$ define $x_n|_b = \sum_{i=b}^{B-1} x_{n,i} 2^i$; namely, $x_n|_b$ is the value that is obtained from x_n by zeroing its b least significant bits, $x_{n,b-1}, \dots, x_{n,0}$. We prove, by induction on $b = B - 1, \dots, 0$, that at the completion of the protocol $x_{n,b} = 1$ iff $x_n|_b = \max\{x_i|_b : i \in [N]\}$. The case $b = 0$ coincides with the statement of the lemma since $x_n|_0 = x_n$.

The statement for the base case $b = B - 1$ is easily proven. There are two cases to consider:

- If $s_b = 1$ then at least one of the inputs has a 1-bit in this column and, therefore, $\max\{x_i|_b : i \in [N]\} = 1$. Hence, after updating the shares in $x_{n,B-1}$ in Line 7, $x_{n,B-1}$ equals the product between $x_{n,B} = 1$ and the original value of $x_{n,B-1}$. That means that $x_{n,B-1}$ remains unchanged in this case and, consequently, $x_{n,B-1} = 1$ iff $x_n|_b = \max\{x_i|_b : i \in [N]\}$.
- If $s_b = 0$ then, because $x_{n,B} = 1$ for all n , all inputs have a 0-bit in this column and, therefore, $\max\{x_i|_b : i \in [N]\} = 0$. Hence, after updating the shares in $x_{n,B-1}$ in Line 7, $x_{n,B-1}$ would equal $x_{n,B} = 1$, for all $n \in [N]$. Indeed, in this case, as all bits in that column are zero, then $x_n|_b = \max\{x_i|_b : i \in [N]\}$ for all $n \in [N]$.

We proceed by induction to prove our claim for $b = B - 2, \dots, 0$. Namely, we assume that $x_n|_{b+1} = \max\{x_i|_{b+1} : i \in [N]\}$, and then prove that $x_n|_b = \max\{x_i|_b : i \in [N]\}$. Here too we distinguish between two cases:

- If $s_b = 0$ then all inputs have a 0-bit in this column or in the preceding column (that was already updated in the previous iteration). We claim that in this case $\arg \max_n \{x_n\}^b = \arg \max_n \{x_n\}^{b+1}$; namely, that all inputs that maximize $x_n\}^{b+1}$ are exactly those that maximize $x_n\}^b$. Indeed, if $i \notin \arg \max_n \{x_n\}^{b+1}$ then clearly $i \notin \arg \max_n \{x_n\}^b$. If, on the other hand, $i \in \arg \max_n \{x_n\}^{b+1}$ then, by the induction hypothesis, $x_{i,b+1} = 1$ and therefore, because $s_b = 0$, $x_{i,b} = 0$. Hence, for all $x_i \in \arg \max_n \{x_n\}^{b+1}$ we have $x_i\}^{b+1} = x_i\}^b$. Therefore, all inputs that maximize $x_n\}^{b+1}$ also maximize $x_n\}^b$. Hence, since the computation in Line 7 updates the shares in $x_{n,b}$ so that it equals $x_{n,b+1}$, the a-posteriori values of the bits $x_{n,b}$ identify the inputs that maximize $x_n\}^b$.
- If $s_b = 1$ then at least one of the inputs that maximized $x_n\}^{b+1}$ has a 1-bit in the b -th column. Therefore, $\arg \max_n \{x_n\}^b$ consists exactly of all inputs in $\arg \max_n \{x_n\}^{b+1}$ that have also a 1-bit in the b -th column. Hence, since in the case $s_b = 1$ the computation in Line 7 updates the shares in $x_{n,b}$ so that it equals the product between $x_{n,b+1}$ and $x_{n,b}$, the a-posteriori values of $x_{n,b}$ identify the set $\arg \max_n \{x_n\}^b$. That completes the proof.

□

Protocol 7: A bit-decomposition-based protocol for MMP

Parameters: N - number of inputs; $B > 0$ - the number of bits in each of the inputs.

Private Inputs: $x_1, \dots, x_N \in [0, 2^B]$; the bit decomposition of x_n is $(x_{n,B-1}, \dots, x_{n,0})$, $n \in [N]$.

- 1 **forall** $n \in [N]$ **do**
- 2 **forall** $b = B - 1, \dots, 0$ **do**
- 3 $[x_{n,b}] \leftarrow (\text{PrivateInput}, n, x_{n,b})$.
- 4 **forall** $b = B - 1, \dots, 0$ **do**
- 5 Compute $[s_b] \leftarrow \bigvee_{n \in [N]} ([x_{n,b+1}] \cdot [x_{n,b}])$.
- 6 **forall** $n \in [N]$ **do**
- 7 $[x_{n,b}] \leftarrow (1 - [s_b]) \cdot [x_{n,b+1}] + [s_b] \cdot [x_{n,b} \cdot x_{n,b+1}]$.
- 8 The parties call $(\text{PublicOutput}, [x_{n,0}])$ for all $n \in [N]$.

Output: The indices $n \in [N]$ for which $x_n = \max_{i \in [N]} x_i$.

We proceed to discuss the protocol's complexity. Each party breaks its input to B binary bits and inputs each of them separately, so that there are in total BN calls to `PrivateInput` (Lines 1-3). Then, there are BN calls to `Multiply` followed by B calls to `Or` (Line 5) and $2BN$ additional calls to `Multiply` (Line 7). Finally, there are N calls to `PublicOutput` (Line 8). In summary,

$$\text{size}(\text{Protocol 7}) = BN \cdot [\text{PrI}] + 3BN \cdot [\text{MuL}] + B \cdot [\text{OR}]_N + N \cdot [\text{PuO}].$$

It implies that

$$\text{preDepth}(\text{Protocol 7}) = \max([\text{PrI}], [\text{MuL}], [\text{OR}]_N, [\text{PuO}]).$$

As for the online depth, let us concentrate on one of the B iterations. All multiplications in Line 5 can be executed in parallel. Once they are completed, the `OR` in Line 5 can be computed. Finally, all $2N$ multiplications in Line 7 can be also parallelized. In summary, we conclude that

$$\text{onDepth}(\text{Protocol 7}) = [\text{PrI}] + B \cdot (2[\text{MuL}] + [\text{OR}]_N) + [\text{PuO}].$$

Example. We conclude our discussion of Protocol 7 by exemplifying its operation. Assume that $N = 4$ and that $x_1 = 11$, $x_2 = 7$, $x_3 = 10$, and $x_4 = 11$. In this case $B = 4$ and the inputs' bits are as follows:

	$b = 3$	$b = 2$	$b = 1$	$b = 0$
$n = 1$	1	0	1	1
$n = 2$	0	1	1	1
$n = 3$	1	0	1	0
$n = 4$	1	0	1	1

We begin with the MSB column, $b = 3$. Here, $s_3 = (1 \cdot 1) \vee (1 \cdot 0) \vee (1 \cdot 1) \vee (1 \cdot 1) = 1$ (since $x_{n,b+1} = x_{n,B} = 1$ in all rows). Hence, the update of bits in that column is $x_{n,3} \leftarrow x_{n,3} \cdot x_{n,4}$ and, because $x_{n,4} = 1$ for all n , the column $b = 3$ remains unchanged.

	$b = 3$	$b = 2$	$b = 1$	$b = 0$
$n = 1$	1	0	1	1
$n = 2$	0	1	1	1
$n = 3$	1	0	1	0
$n = 4$	1	0	1	1

Next, we deal with $b = 2$. Here $s_2 = (1 \cdot 0) \vee (0 \cdot 1) \vee (1 \cdot 0) \vee (1 \cdot 0) = 0$. Hence, the update rule is $x_{n,2} \leftarrow x_{n,3}$. Therefore, when completing this iteration we end up with the following table of bits:

	$b = 3$	$b = 2$	$b = 1$	$b = 0$
$n = 1$	1	1	1	1
$n = 2$	0	0	1	1
$n = 3$	1	1	1	0
$n = 4$	1	1	1	1

Moving to $b = 1$ we see that $s_1 = 1$ and, therefore, the update rule is $x_{n,1} \leftarrow x_{n,1} \cdot x_{n,2}$. As a result, we end up with the following table:

	$b = 3$	$b = 2$	$b = 1$	$b = 0$
$n = 1$	1	1	1	1
$n = 2$	0	0	0	1
$n = 3$	1	1	1	0
$n = 4$	1	1	1	1

Finally, we reach the LSB, $b = 0$. Here too $s_0 = 1$. Hence, after the update $x_{n,0} \leftarrow x_{n,0} \cdot x_{n,1}$ we get

	$b = 3$	$b = 2$	$b = 1$	$b = 0$
$n = 1$	1	1	1	1
$n = 2$	0	0	0	0
$n = 3$	1	1	1	0
$n = 4$	1	1	1	1

Indeed, the bits in the column $b = 0$ identify the maximal inputs in rows $n = 1$ and $n = 4$.

5.2 Reducing depth by using larger bases

The number of iterations in the main loop in Protocol 7 directly determines its depth, since each iteration is computed by a constant-depth sub-protocol. The number of iterations equals the number B of bits in the representation of the inputs. Protocol 7 can be modified so that a larger base 2^d , $d > 1$, is used. Such a modification yields a smaller number of iterations $- B/d$ instead of B (for the sake of

simplicity we assume that $d|B$). The generalization is described in Protocol 8, which uses a constant-depth sub-protocol for computing maxima (the same technique as used in Protocol 5).

Protocol 8: A digit-decomposition-based protocol for MMP

Parameters: N - number of inputs in $[0, 2^B)$; $d \in \mathbb{N}$ s.t. $d|B$;
 $Q := 2^d$; $D := B/d$ is the number of base- 2^d
 digits in each of the inputs.

Private Inputs: $x_1, \dots, x_N \in [0, 2^B)$; the digit
 decomposition of $x_n, n \in [N]$, is
 $(x_{n,D-1}, \dots, x_{n,0})$, where $x_{n,\ell} \in [0, Q)$,
 $0 \leq \ell \leq D-1$ (that is, $x_n = \sum_{\ell=0}^{D-1} x_{n,\ell} \cdot Q^\ell$).

```

1 forall n in [N] do
2   forall 0 ≤ ℓ ≤ D - 1 do
3     forall i in [Q - 1] do
4       | [yn,ℓ,i] ← (PrivatInput, n, yn,ℓ,i), where
         | yn,ℓ,i = 1xn,ℓ ≥ i.
5 Initialize [rn,D] ← 1 for all n in [N].
6 forall ℓ = D - 1, ..., 0 do
7   forall n in [N] and i in [Q - 1] do
8     | [y'n,ℓ,i] ← [yn,ℓ,i] · [rn,ℓ+1]
9     ([Mℓ,Q-1], ..., [Mℓ,1]) ←
       MaxMonotone(N, Q, {[y'n,ℓ,i]}n∈[N],i∈[Q-1]).
10  forall n in [N] do
11    | [1Mℓ=xn,ℓ] =
       EqualBits(Q, {[y'n,ℓ,i]}i∈[Q-1], {Mℓ,i}i∈[Q-1]).
12    | [rn,ℓ] ← [rn,ℓ+1] · [1Mℓ=xn,ℓ].
13 The parties call (PublicOutput, [rn,0]) for all n in [N].
Output: The indices n in [N] for which
         xn = M := max{xi : i in [N]}.

```

In Protocol 8 the parties decompose their inputs to digits in a 2^d -base. Thus, if $x_n \in [0, 2^B)$, the number of 2^d -digits in it is $D = B/d$, where the digit in the ℓ -th position is denoted $x_{n,\ell}$, $0 \leq \ell < D$. Each such digit is an integer in the range $[0, Q-1]$ where $Q = 2^d$. Then, each digit is provided as an input to the protocol by its monotone representation; i.e., the digit $x_{n,\ell}$ is represented by $Q-1$ bits – $y_{n,\ell,1}, \dots, y_{n,\ell,Q-1}$ – such that $y_{n,\ell,i} = 1$ iff $x_{n,\ell} \geq i$. Therefore, the representation is monotonically decreasing: $y_{n,\ell,i} \geq y_{n,\ell,i+1}$ for every $1 \leq i < Q-1$. The input of all $Q-1$ bits in the monotone representation of all D digits in all N inputs is done in Lines 1-4.

The variables $r_{n,\ell}$, $n \in [N]$, $\ell = D, \dots, 0$, are binary flags that equal 1 iff the n -th input is still a candidate to be the maximal input after scanning all inputs' digits from the $(D-1)$ -th digit down to the ℓ -th digit. Since initially, before the scan starts, all inputs are candidates, those variables are initialized to $r_{n,D} = 1$ for all $n \in [N]$ (Line 5).

The main loop (Lines 6-12) scans the D digits of all inputs, from the most significant one ($\ell = D-1$) to the least significant one ($\ell = 0$). First, the parties multiply the monotone bit representation of the current bit in each of the inputs with the bit that indicates whether that input is still a potential candidate to be the maximum (Lines 7-8). Then, the monotone representations of the current digit in each of the relevant inputs are used to compute the monotone

representation of the maximal digit in that position using the sub-protocol MaxMonotone (Line 9). Then, we check the digits of all inputs against the computed maximum (Lines 10-11) by calling the sub-protocol EqualBits. Subsequently (Line 12), we update $r_{n,\ell}$ to be 1 iff that input is still a candidate at this point ($r_{n,\ell+1} = 1$) and the current digit is maximal ($M_\ell = x_{n,\ell}$). At the conclusion of this loop, the protocol outputs the bits $r_{n,0}$, $n \in [N]$, that identify the maximal inputs (Line 13). The proof of this protocol's correctness is very similar to that of Protocol 7, see Lemma 4.

We now turn to describe the sub-protocols that Protocol 8 invokes. Sub-protocol 9 computes the monotone representation of the maximum of inputs that are also given through their monotone representation. It is called by Protocol 8 in Line 9. Sub-protocol 10, namely EqualBits($Q, \{[x_i]\}_{i \in [Q-1]}, \{[x'_i]\}_{i \in [Q-1]}$), is given handles to two bit-vectors and it outputs 1 iff they are equal in every position. We rely here on the straightforward identity $\bigwedge_{i \in [Q-1]} (1_{x_i=x'_i}) = 1 - (\bigvee_{i \in [Q-1]} (x_i - x'_i)^2)$. This sub-protocol is called in Line 11 of Protocol 8.

Sub-protocol 9: MaxMonotone($N, Q, \{[y_{n,i}]\}_{n \in [N], i \in [Q-1]}$)

Parameters: N - number of inputs; Q - an upper bound on the inputs (each input is in $[0, Q-1]$).

Inputs: Handles $[y_{n,i}]$, $n \in [N]$, $i \in [Q-1]$, such that $y_{n,i} \in \{0, 1\}$ and $y_{n,i} \geq y_{n,i+1}$ for all $i \in [Q-2]$ and $n \in [N]$.

```
1 forall i in [Q - 1] do
```

```
2 | [Mi] ←  $\bigvee_{n \in [N]} [y_{n,i}]$ .
```

Output: The handles ($[M_{Q-1}], \dots, [M_1]$) to the monotone representation of

$$M := \max_{n \in [N]} \{\max_{i \in [Q-1]} \{y_{n,i}\}\}.$$

Sub-protocol 10: EqualBits($Q, \{[x_i]\}_{i \in [Q-1]}, \{[x'_i]\}_{i \in [Q-1]}$)

Parameters: Q - determines the length of the two bit vectors.

Inputs: Handles to two bit vectors, $\{[x_i]\}_{i \in [Q-1]}$ and $\{[x'_i]\}_{i \in [Q-1]}$.

```
1 [1x=x'] ← 1 - ( $\bigvee_{i \in [Q-1]} ([x_i] - [x'_i])^2$ ).
```

Output: The handle [1_{x=x'}].

To analyze the complexity of Protocol 9 we start by analyzing the complexity of the sub-protocols. The costs of MaxMonotone are:

$$\begin{aligned} \text{size(Protocol 9)} &= (Q-1)[\text{OR}]_N, \\ \text{preDepth(Protocol 9)} &= [\text{OR}]_N, \\ \text{onDepth(Protocol 9)} &= [\text{OR}]_N. \end{aligned}$$

The costs of EqualBits are:

$$\begin{aligned} \text{size(Protocol 10)} &= (Q-1)[\text{MUL}] + [\text{OR}]_{Q-1}, \\ \text{preDepth(Protocol 10)} &= \max([\text{MUL}], [\text{OR}]_{Q-1}), \\ \text{onDepth(Protocol 10)} &= [\text{MUL}] + [\text{OR}]_{Q-1}. \end{aligned}$$

Hence, Protocol 8's costs are:

$$\begin{aligned} \text{size(Protocol 8)} &= ND(Q-1)[\text{PrI}] + ND(Q-1)[\text{MuL}] + \\ &D[\text{MaxMonotone}] + ND([\text{EqualBits}] + [\text{MuL}]) + N[\text{PuO}] = \\ &ND(Q-1)[\text{PrI}] + ND(Q-1)[\text{MuL}] + D(Q-1)[\text{OR}]_N + \\ &ND((Q-1)[\text{MuL}] + [\text{OR}]_{Q-1} + [\text{MuL}]) + N[\text{PuO}] = \\ &ND(Q-1)[\text{PrI}] + ND(2Q-1)[\text{MuL}] + ND[\text{OR}]_{Q-1} + \\ &D(Q-1)[\text{OR}]_N + N[\text{PuO}] \\ \text{preDepth(Protocol 8)} &= \\ &\max([\text{PrI}], [\text{MuL}], [\text{OR}]_N, [\text{OR}]_{Q-1}, [\text{PuO}]), \\ \text{onDepth(Protocol 8)} &= [\text{PrI}] + \\ &D([\text{MuL}] + [\text{MaxMonotone}] + [\text{EqualBits}] + [\text{MuL}]) + \\ &[\text{PuO}] = [\text{PrI}] + D(3[\text{MuL}] + [\text{OR}]_N + [\text{OR}]_{Q-1}) + [\text{PuO}]. \end{aligned}$$

6 A PROTOCOL FOR MMP INSPIRED BY [1]

The study of Aggarwal et al. [1] presented a secure protocol for finding the k -th ranked element in the union of private datasets. It can be used to solve MaxP, and subsequently MMP, if we view the input to those problems as N private singleton datasets and take $k = N$. We proceed to describe such a protocol, for the sake of comparing it with the previous protocols that were presented in this study.

Protocol 11 computes the maximum by implementing a privacy-preserving binary search over the interval $[0, 2^B - 1]$, where $2^B - 1$ is an upper bound on all inputs (Lines 1-15). The lower and upper bounds of the search interval, denoted a and b , are initialized in Line 1. A binary flag f is set to zero in Line 2; it will be set to 1 once the maximum is found. The binary search loop is given in Lines 3-15. The contemplated value of the maximum, M , is set to the middle of the current search interval (Line 4). Then, each of the parties set two local flags, ℓ_n and g_n , that indicate whether its input is larger or smaller than M (Lines 5-7). Subsequently, all determine jointly whether there is at least one input larger than M ($u = 1$) or not ($u = 0$), and whether all inputs are smaller than M ($v = 1$) or not ($v = 0$) (Line 8). This is done by computing the OR (AND) of the private bits ℓ_n (g_n).

If $x_n \leq M$ for all n ($u = 0$) and there exists at least one input x_j such that $x_j \geq M$ ($v = 0$) then M is the sought-after maximum. In that case f is set to 1 so that the loop terminates (Lines 9-10). Otherwise, if $u = 1$ then there is at least one input larger than M so that the sought maximum is at least $M + 1$. In that case we update the search lower bound to that value (Line 11-12). Otherwise, if $v = 1$ then all inputs are smaller than M and hence the search upper bound is updated to $M - 1$ (Lines 13-14).

In the concluding stage, the protocol finds all private inputs that equal the maximum (Lines 16-17) and outputs them.

7 CONCRETE EVALUATION

The goal of this section is to arm the reader with a methodology of measuring and comparing the protocols presented in this paper for a given implementation of the ABB functionality in a given setting. To demonstrate the methodology, we pick the statistically secure implementation of Damgård and Nielsen [11], in the setting of semi-honest adversary, who statically corrupts a minority of the parties (also known as the honest majority setting). In Section 7.1 we present an implementation of the 'high-level' interfaces based on the implementation of the 'low-level' ones, as done by Damgård

Protocol 11: Solving MMP by binary search [1]

Private inputs: P_n has $x_n \in [0, 2^B - 1]$, $n \in [N]$.

- 1 Set $a \leftarrow 0$ and $b \leftarrow 2^{B-1}$
- 2 Set $f \leftarrow 0$.
- 3 **repeat**
- 4 Set $M \leftarrow \lceil (a+b)/2 \rceil$
- 5 **forall** $n \in [N]$ **do**
- 6 P_n sets $\ell_n \leftarrow 1$ if $x_n > M$ and $\ell_n \leftarrow 0$ otherwise.
- 7 P_n sets $g_n \leftarrow 1$ if $x_n < M$ and $g_n \leftarrow 0$ otherwise.
- 8 The parties set $u \leftarrow \bigvee_n \ell_n$ and $v \leftarrow \bigwedge_n g_n$.
- 9 **if** $u = 0$ and $v = 0$ **then**
- 10 Set $f \leftarrow 1$.
- 11 **else if** $u = 1$ **then**
- 12 Set $a \leftarrow M + 1$.
- 13 **else if** $v = 1$ **then**
- 14 Set $b \leftarrow M - 1$.
- 15 **until** $f = 1$.
- 16 **forall** $n \in [N]$ **do**
- 17 The parties compute $1_{x_n=M}$.

Output: The indices $n \in [N]$ for which $x_n = M := \max\{x_i : i \in [N]\}$.

and Nielsen [11]. Then we formulate the costs of all interfaces based on the specific parameters of the Multiple Millionaires' Problem in Section 7.2. These parameters consist of the number of parties, number of inputs and the bit-length of the inputs. Finally, in Section 7.3 we compare the performance of the protocol for a set of specific parameters.

7.1 Interfaces

Here, we break down the high-level interfaces (like Equality and OR) into their cost in terms of more primitive interfaces (like PrivateInput and Multiply), and then derive their concrete costs, for a specific setting that we selected. Such a detailed analysis is necessary in order to enable a comparison between protocols and in order to get a more precise estimate of the costs of a given protocol. However, we stress that such a breakdown analysis might be different for different settings and, hence, it should be carried out independently for each such setting.

The cost analysis of interfaces [PrI], [PuO], [RNd], [MuL] is based on the passively secure MPC protocol of Damgård and Nielsen [11]; the cost analysis of interface [Inv] is trivial and given in [30, Appendix A.2.1]; and the cost of interfaces [COM] and [OR]_m is as derived in [17] and [26], respectively.

We examine a setting where the ABB is implemented by a protocol that guarantees statistical security in the presence of a semi-honest adversary. We choose a semi-honest adversary in order to avoid noisy costs that are typically added to ensure input consistency when the adversary is malicious. These costs might distract us from a clean and objective comparison between the protocols.

As stated in the Introduction, we distinguish between the number of inputs, N , and the number of MPC parties, K . We utilize this distinction to get more accurate measures for the size of the

presented protocols, being the communication complexity per MPC party (the number of field elements it sends/receives).

The setting that we picked allows us to use the concrete costs of Or and Equality (see [30, Appendix A.2]). In the following we present the final costs, after integrating the concrete costs from [30, Appendix A.2]. Let us first provide concrete costs to some of the interfaces in Functionality 1:

- **PrivateInput:** Submitting a private input boils down to a simple a secret sharing, in which the party who acts as the dealer sends 1 field element to every other party, in 1 round. Hence, $\text{size}([\text{PRI}]) = 1$, $\text{preDepth}([\text{PRI}]) = 0$, $\text{onDepth}([\text{PRI}]) = 1$.

- **PublicOutput** is implemented by opening the secret to a single designated party, who then announces the result to all other parties. This translates into a communication of 2 field elements per party, in 2 rounds. The number of communication rounds may be reduced, at the expense of increasing the communication cost, in the following manner: each party sends its share in the output to all other parties; that entails sending K field elements in 1 round. We use the latter protocol as the basis to our calculations, hence $\text{size}([\text{PUO}]) = K$, $\text{preDepth}([\text{PUO}]) = 0$, $\text{onDepth}([\text{PUO}]) = 1$.

- **Random:** Generating a single secret random field element and distributing shares in it entails communicating 1 field element (per party) in 1 round. Hence, $\text{size}([\text{RND}]) = 1$, $\text{preDepth}([\text{RND}]) = 1$, $\text{onDepth}([\text{RND}]) = 0$.

- **(Multiply, $[a], [b]$)** is implemented by the DoubleRandom technique presented in [11]. The DoubleRandom procedure outputs two handles, $[r]$ and $[R]$, to the same random field element, where $[r]$ is a sharing of degree t and $[R]$ is a sharing of degree $2t$, and t is an upper bound on the number of corrupted parties. The parties locally compute $[C] = [a] \cdot [b]$, in order to get a sharing $[C]$ of the product $a \cdot b$ which is of degree $2t$. Then they perform PublicOutput on $c' = [C] - [R]$ and locally compute $[c] = [a \cdot b] = c' + [r]$. DoubleRandom requires twice more communication than Random and the same round complexity (which is one round). Overall, the communication cost is $2 + K$ field elements per party, and 2 rounds, where the DoubleRandom's round occurs in the preprocessing stage, whereas the PublicOutput's round occurs in the online phase. Then,

$\text{size}([\text{MUL}]) = 2 + K$, $\text{preDepth}([\text{MUL}]) = 1$, $\text{onDepth}([\text{MUL}]) = 1$.

- **(Inverse, $[a]$)** can be computed by invoking Random, Multiplly and PublicOutput (see [30, Appendix A.2.1]). The costs of that computation are as follows:

$\text{size}([\text{INV}]) = 3 + 2K$, $\text{preDepth}([\text{INV}]) = 1$, $\text{onDepth}([\text{INV}]) = 2$.

- **(Equality, $[a], [b]$)** too (like Inverse) can be computed by invoking Random, Multiplly and PublicOutput (see [30, Appendix A.2.2]). Hence, its costs are:

$\text{size}([\text{EQV}]) = 3 + 2K$, $\text{preDepth}([\text{EQV}]) = 1$, $\text{onDepth}([\text{EQV}]) = 2$.

- **(Compare, $[a], [b]$):** We use the comparison protocol of Goss and Jiang [17]. The costs of that protocol are:

$\text{size}([\text{COM}]) = 2B + 5 \log B + 23$,

$\text{preDepth}([\text{COM}]) = 0$, $\text{onDepth}([\text{COM}]) = 4$.

- **(OR_m, $[a_1], \dots, [a_m]$):** Nishide and Ohta [26] designed a protocol for computing $[\bigvee_{i \in [m]} a_i]$ in constant depth. (In [30, Appendix A.2.3] we provide a description of their solution.) The costs of their protocol are:

$\text{size}([\text{OR}]_m) = 2m \cdot [\text{RND}] + (3m - 1) \cdot [\text{MUL}] + 2m \cdot [\text{PUO}] = 2m + (3m - 1)(2 + K) + 2mK = 5NK + 8m - K - 2$.

In the preprocessing stage there are calls to Random, followed by Multiply and then PublicOutput, which are dependent on each other. Therefore, the preprocessing depth equals the sum of the *total* depths of these interfaces, which is 4, because $\text{depth}(\text{Random}) = \text{depth}(\text{PublicOutput}) = 1$ and $\text{depth}(\text{Multiply}) = 2$. However, note that Multiply consists of an invocation of Random on its own, which can be executed in the first round (rather than in the second), and that saves one round. Overall, there are 3 rounds in the preprocessing stage:

$\text{preDepth}([\text{OR}]_m) = [\text{RND}] + [\text{MUL}] + [\text{PUO}] = 3$.

As for the online depth,

$\text{onDepth}([\text{OR}]_m) = 3 \cdot [\text{MUL}] + [\text{PUO}] = 4$.

7.2 Protocols' costs

Now we may proceed to derive the costs of the protocols. We do so by taking the costs (size and depth) of the protocols as we analyzed them in Sections 3–6 and plugging in them the costs of the interfaces that they invoke, as given in Section 7.1. The parameters that determine those costs are: N - number of inputs, K - number of MPC parties, and B - number of bits for representing all inputs. Due to page limitation we omit the detailed computation of those costs and refer the readers to [30, Appendix B]. In Table 1 we provide a summary of those costs (where the size is shown only asymptotically, while the preprocessing and online depths are shown exactly).

	size	preDepth	onDepth
Protocol 1	$O(NB + NK)$	1	$5 \cdot \lceil \log N \rceil + 4$
Protocol 4	$O(N'NB + N'NK)$	3	$3 + 12 \lceil \log_{N'} N \rceil$
Protocol 6	$(2^B NK)$	3	9
Protocol 7	$O(NBK)$	3	$2 + 6B$
Protocol 8	$O(\frac{NBK2^d}{d})$	3	$2 + 11B/d$
Protocol 11	$O(NB^2 + NBK)$	3	$8B + 4$

Table 1: Analysis of the size, preprocessing depth and online depth of the protocols presented in this paper, according to Definitions 2-3 (see Section 2.2), when instantiated in a setting of honest majority, semi-honest adversary and statistical security (See Section 7 for details). The size is given asymptotically, while the preprocessing depth and online depth are given exactly. The parameters are: K is the number of parties; N is the number of inputs; each input is drawn from the domain $[0, 2^B - 1]$; $N' \in [2, N]$ is the degree of the general comparison tree in Protocol 4; d is the bit-length of digits in Protocol 8.




Figure 1: The size and depth the protocols presented in this paper.

7.3 Concrete costs in specific scenarios

We present the concrete cost for two scenarios: in the first (Figure 1, top), the inputs x_n are drawn from a small domain, $0 \leq x_n < 256$ and so $B = 8$; and in the second (Figure 1, bottom), the inputs are drawn from a large domain, $0 \leq x_n < 2^{32}$, and so $B = 32$. In the figure, the protocols are distinguished by color and line style: Protocol 1 is presented by a solid black line; Protocol 4 has three versions, each with a different setting of N' , presented by green lines (solid, dashed and dotted); Protocol 6 is presented by a solid blue line; Protocol 7 is presented by a pink line; Protocol 8 is presented by orange lines, where the line style (solid, dashed or dotted) identify the setting of d , the digit length; and Protocol 11 is presented by a solid cyan line. The two plots in the figure present the trend of the costs as the number of inputs, N , grows. We evaluate the costs for $N \in \{2^8, 2^{16}, 2^{24}, 2^{32}\}$, and each setting is presented by a

different marker (see the legend on the left side of each of the two plots). For other scenarios, the reader may access the notebook with the source code that we used to calculate these costs and change the parameters accordingly, see https://colab.research.google.com/drive/1yjEt0lhFteSIUb06ipjxCzKSnnE5PTfA?usp=drive_link.

When latency is at premium (e.g., when a system invokes MMP repetitively) one would prefer using a low depth protocol. The comparison in Figure 1 suggests that when the inputs are taken from a small domain (e.g., $B = 8$) then the monotone-representation-based protocol (Protocol 6) performs best. If, however, the inputs are taken from larger domains (e.g., $B = 32$), Protocol 6 becomes impractical and then one should consider Protocol 8 or Protocol 4, with a suitable choice of their parameters (d and N' , respectively). On the other hand, when bandwidth is at premium, the “naïve” binary-tree-based protocol (Protocol 1) performs better than others.

8 RELATED WORK

While many papers considered the 2-party Millionaires' Problem, its multi-party generalization, MMP, is introduced and studied herein for the first time (to the best of our knowledge). We review here some studies that addressed problems similar to MMP and MaxP.

Liu et. al. [23] considered a problem that is related to MaxP. They assumed that all inputs in $I := \{x_n : n \in [N]\}$ are distinct non-negative integers that are bounded by a small integer Q . They devised a protocol that outputs the set I completely, while perfectly hiding the mapping between I and the set of parties $\{P_n : n \in [N]\}$. While such a solution reveals the output of MaxP, it reveals far more than what is desired in MaxP and MMP. In addition, the complexity of their protocol scales linearly with Q (which is equivalent to our 2^B), what limits significantly its scalability.

The studies [28, 29] concentrate on the sequencing problem: the setting in that problem is as in MMP and MaxP but the required output is a set of ranks $\{r_m : m \in [N]\}$, where $r_m := |\{n \in [N] : x_n > x_m\}|$, $m \in [N]$. Hence, all inputs x_m for which $r_m = 0$ are the maximum. Here, as opposed to the above mentioned study of Liu et. al. [23], the values of the inputs remain hidden. But the output is still far more elaborated than what is desired in MMP and MaxP.

Aggarwal et al. [1] considered a setting with two or more parties possessing large confidential datasets and designed secure protocols for computing the k -th ranked element of the union of the datasets. Their protocol template FIND-RANKED-ELEMENT-MULTIPARTY can be translated into a protocol for solving MaxP. We presented that protocol as Protocol 11 in Section 6. The solution to MaxP as issued by Protocol 11 can be securely converted to a solution of MMP (as we do in Lines 12-13 in Protocol 1). In Section 6 we also analyzed the protocol's size and depth, and included it in our comparison in Table 1 and in Figure 1 in Section 7.3.

David et al. [12] and Mohassel et. al. [24] solved the MMP problem as a building block for constructing privacy-preserving machine learning protocols. The approach taken by these works is similar to the binary-tree-based protocol (Protocol 1) and has the same costs. Namely, they solve MMP by reducing it to a sequence of (two party) MPs. In [12], they also mentioned a 'flattened tree' approach, in which the tree's height is 1. This means that each leaf is compared to every other leaf at the same round. This approach is captured by our general-tree-based protocol (Protocol 4) by setting $N' = N$.

Recall that our bit-decomposition and digit-decomposition based protocols (Protocols 7 and 8, respectively) have depth *linear* in the number of digits, D ($1 \leq D \leq B$). On the other hand, protocols that solve the original (two-party) Millionaires' problem have depth that is only *logarithmic* in D , see [9, 14, 27]. These protocols use the following recurrence relation: to compare two numbers x and y of B bits each, we can split them to the four integers $[x_{\text{high}}, x_{\text{low}}]$ and $[y_{\text{high}}, y_{\text{low}}]$, where x_{high} (resp. y_{high}) is of B_{high} bits, x_{low} (resp. y_{low}) is of $B_{\text{low}} := B - B_{\text{high}}$ bits, and $x = x_{\text{high}} \cdot 2^{B_{\text{low}}} + x_{\text{low}}$ (resp. $y = y_{\text{high}} \cdot 2^{B_{\text{low}}} + y_{\text{low}}$). This splitting enables computing the comparison bit $1_{x < y}$ through the equation

$$1_{x < y} = 1_{x_{\text{high}} < y_{\text{high}}} \vee (1_{x_{\text{high}} = y_{\text{high}}} \wedge 1_{x_{\text{low}} < y_{\text{low}}}), \quad (1)$$

which leads to a binary tree of comparisons of depth $\log D$. Let us try to extend the above recurrence relation to MMP, and let us consider the case $N = 3$. Letting x, y, z be the three private inputs,

each consisting of B bits, and each is split into lower and higher parts, as described above, then the natural extension of Eq. (1) would be as follows:

$$1_{x = \max(x, y, z)} = 1_{x_{\text{high}} = \max(x_{\text{high}}, y_{\text{high}}, z_{\text{high}})} \\ \vee (1_{x_{\text{high}} = y_{\text{high}} = z_{\text{high}}} \wedge 1_{x_{\text{low}} = \max(x_{\text{low}}, y_{\text{low}}, z_{\text{low}})}).$$

However, it is incorrect. For instance, if $x_{\text{high}} = y_{\text{high}} > z_{\text{high}}$ and $y_{\text{low}} > x_{\text{low}}$, then a protocol based on that relation would output both x and y as the maximal values, whereas the correct output is only y . Therefore, it is not clear whether the two-party recurrence relation can be extended to the multi-party setting.

9 CONCLUSION

We studied here two fundamental MPC problems — MMP and MaxP. Those problems are natural extensions of Yao's classical Millionaires' Problem [31]. As applications of privacy-preserving computation are more and more commonly implemented in industrial systems, MMP and MaxP become important building blocks in privacy-preserving statistics, machine learning, auctions and other domains.

While some prior studies considered problems related to MMP [1, 23, 28, 29], and others did solve MMP by reducing it to a sequence of MPs [12, 24], it appears that ours is the first study that introduces this fundamental MPC problem and offers dedicated solutions, of different approaches, and compares between them in order to illustrate the tradeoff between the size and the depth of the corresponding protocols.

A prominent advantage of our novel protocols is their simplicity. As they solve fundamental problems that are essential building blocks in important applications, our study of solutions to those basic problems, and the comparison between them, will serve well future researchers of MPC and practitioners of secure distributed computing.

10 ACKNOWLEDGEMENTS

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

REFERENCES

- [1] Gagan Aggarwal, Nina Mishra, and Benny Pinkas. 2010. Secure Computation of the Median (and Other Elements of Specified Ranks). *J. Cryptol.* 23 (2010), 373–401.
- [2] Ramiro Alvarez and Mehrdad Nojoumian. 2020. Comprehensive survey on privacy-preserving protocols for sealed-bid auctions. *Comput. Secur.* 88 (2020). <https://doi.org/10.1016/j.cose.2019.03.023>
- [3] Ian F. Blake and Vladimir Kolesnikov. 2004. Strong Conditional Oblivious Transfer and Computing on Intervals. In *ASIACRYPT*. 515–529.
- [4] Marina Blanton and Paolo Gasti. 2011. Secure and Efficient Protocols for Iris and Fingerprint Identification. In *ESORICS*. 190–209.
- [5] Paul Bunn and Rafail Ostrovsky. 2007. Secure Two-Party k-Means Clustering. *IACR Cryptol. ePrint Arch.* (2007), 231. <http://eprint.iacr.org/2007/231>
- [6] Ran Canetti. 2000. Security and Composition of Multiparty Cryptographic Protocols. *J. Cryptol.* 13, 1 (2000), 143–202. <https://doi.org/10.1007/s001459910006>
- [7] Yao-Jen Chang, Chia-Wei Tsai, and Tzonelih Hwang. 2013. Multi-user private comparison protocol using GHZ class states. *Quantum Inf. Process.* 12 (2013), 1077–1088.
- [8] Hao Chen, Ilaria Chillotti, Yihe Dong, Oxana Poburinnaya, Ilya P. Razenshteyn, and M. Sadegh Riazi. 2020. SANNS: Scaling Up Secure Approximate k-Nearest Neighbors Search. In *USENIX*. 2111–2128.
- [9] Geoffroy Couteau. 2018. New Protocols for Secure Equality Test and Comparison. In *ACNS*. 303–320.
- [10] Ivan Damgård and Jesper Buus Nielsen. 2003. Universally Composable Efficient Multiparty Computation from Threshold Homomorphic Encryption. In *CRYPTO*. 247–264.
- [11] Ivan Damgård and Jesper Buus Nielsen. 2007. Scalable and Unconditionally Secure Multiparty Computation. In *CRYPTO*. 572–590.
- [12] Bernardo Machado David, Rafael Dowsley, Raj S. Katti, and Anderson C. A. Nascimento. 2015. Efficient Unconditionally Secure Comparison and Privacy Preserving Machine Learning Classification Protocols. In *ProvSec*. 354–367.
- [13] Marc Fischlin. 2001. A Cost-Effective Pay-Per-Multiplication Comparison Method for Millionaires. In *Topics in Cryptology - CT-RSA*. 457–472.
- [14] Juan A. Garay, Berry Schoenmakers, and José Villegas. 2007. Practical and Secure Solutions for Integer Comparison. In *PKC*. 330–342.
- [15] Daniel Genkin, Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and Eran Tromer. 2014. Circuits resilient to additive attacks with applications to secure computation. In *STOC 2014*, David B. Shmoys (Ed.). ACM, 495–504.
- [16] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or a Completeness Theorem for Protocols with Honest Majority. In *STOC*. 218–229.
- [17] Ken Goss and Wei Jiang. 2018. Efficient and Constant-Rounds Secure Comparison through Dynamic Groups and Asymmetric Computations. *IACR Cryptol. ePrint Arch.* (2018), 179. <http://eprint.iacr.org/2018/179>
- [18] Dima Grigoriev, Laszlo B. Kish, and Vladimir Shpilrain. 2017. Yao’s Millionaires’ Problem and Public-Key Encryption Without Computational Assumptions. *Int. J. Found. Comput. Sci.* 28 (2017), 379–390.
- [19] Yan Huang, Lior Malka, David Evans, and Jonathan Katz. 2011. Efficient Privacy-Preserving Biometric Identification. In *NDSS*.
- [20] Ioannis Ioannidis and Ananth Grama. 2003. An Efficient Protocol for Yao’s Millionaires’ Problem. In *Hawaii International Conference on System Sciences (HICSS)*. 205.
- [21] Kimmo Järvinen, Helena Leppäkoski, Elena Simona Lohan, Philipp Richter, Thomas Schneider, Oleksandr Tkachenko, and Zheng Yang. 2019. PILOT: Practical Privacy-Preserving Indoor Localization Using Outsourcing. In *European Symposium on Security and Privacy, EuroS&P*. 448–463.
- [22] Hsiao-Ying Lin and Wen-Guey Tzeng. 2005. An Efficient Solution to the Millionaires’ Problem Based on Homomorphic Encryption. In *ACNS*. 456–466.
- [23] Xin Liu, Shundong Li, Xiubo Chen, Gang Xu, Xiaolin Zhang, and Yong Zhou. 2017. Efficient Solutions to Two-Party and Multiparty Millionaires’ Problem. *Secur. Commun. Networks* 2017 (2017), 5207386:1–5207386:11. <https://doi.org/10.1155/2017/5207386>
- [24] Payman Mohassel, Mike Rosulek, and Ni Trieu. 2020. Practical Privacy-Preserving K-means Clustering. *Proc. Priv. Enhancing Technol.* 2020 (2020), 414–433.
- [25] Moni Naor, Benny Pinkas, and Reuban Sumner. 1999. Privacy preserving auctions and mechanism design. In *EC-99*. 129–139.
- [26] Takashi Nishide and Kazuo Ohta. 2007. Multiparty Computation for Interval, Equality, and Comparison Without Bit-Decomposition Protocol. In *PKC*. 343–360.
- [27] Deevashwer Rathee, Mayank Rathee, Nishant Kumar, Nishanth Chandran, Divya Gupta, Aseem Rastogi, and Rahul Sharma. 2020. CryptFlow2: Practical 2-Party Secure Inference. In *CCS*. 325–342.
- [28] Yi Sun, Qiaoyan Wen, Yudong Zhang, Hua Zhang, and Zhengping Jin. 2013. Efficient Secure Multiparty Computation Protocol for Sequencing Problem over Insecure Channel. *Artificial Intelligence and Its Applications* (2013). Article ID 172718.
- [29] ChunMing Tang, GuiHua Shi, and ZhengAn Yao. 2011. Secure multi-party computation protocol for sequencing problem. *Sci. China Inf. Sci.* 54 (2011), 1654–1662.
- [30] Tamir Tassa and Avishay Yanai. 2024. The Multiple Millionaires’ Problem. *IACR Cryptol. ePrint Arch.* (2024), 5. <https://eprint.iacr.org/2024/005>
- [31] Andrew C. Yao. 1982. Protocols for secure computation. In *FOCS*. 160–164.