

# Prime Match: A Privacy-Preserving Inventory Matching System

*Antigoni Polychroniadou*      *Gilad Asharov*<sup>1</sup>      *Benjamin Diamond*<sup>2</sup>      *Tucker Balch*  
*Hans Buehler*<sup>2</sup>      *Richard Hua*      *Suwen Gu*      *Greg Gimler*<sup>2</sup>      *Manuela Veloso*

*J.P. Morgan*

## Abstract

Inventory matching is a standard mechanism for trading financial stocks by which buyers and sellers can be paired. In the financial world, banks often undertake the task of finding such matches between their clients. The related stocks can be traded without adversely impacting the market price for either client. If matches between clients are found, the bank can offer the trade at advantageous rates. If no match is found, the parties have to buy or sell the stock in the public market, which introduces additional costs.

A problem with the process as it is presently conducted is that the involved parties must share their order to buy or sell a particular stock, along with the intended quantity (number of shares), to the bank. Clients worry that if this information were to “leak” somehow, then other market participants would become aware of their intentions and thus cause the price to move adversely against them before their transaction finalizes.

We provide a solution, Prime Match, that enables clients to match their orders efficiently with reduced market impact while maintaining privacy. In the case where there are no matches, no information is revealed. Our main cryptographic innovation is a two-round secure *linear* comparison protocol for computing the minimum between two quantities without preprocessing and with malicious security, which can be of independent interest. We report benchmarks of our Prime Match system, which runs in production and is adopted by a large bank in the US – J.P. Morgan. The system is designed utilizing a star topology network, which provides clients with a centralized node (the bank) as an alternative to the idealized assumption of point-to-point connections, which would be impractical and undesired for the clients to implement in reality.

Prime Match is the first secure multiparty computation solution running live in the traditional financial world.

<sup>1</sup>Department of Computer Science, Bar-Ilan University. Work initiated while at J.P. Morgan AI Research. Editorial work while at Bar-Ilan University, supported by J.P. Morgan Faculty Research Award.

<sup>2</sup>Authors contributed while employed at J.P. Morgan. Buehler moved to XTX Markets, Diamond to Coinbase, and Gimler to Meta.

## 1 Introduction

An *axe* is an interest in a particular stock that an investment firm wishes to buy or sell. Banks and brokerages provide their clients with a matching service, referred to as “*axe matching*”. When a bank finds two clients interested in the same stock but with opposite directions (one is interested in buying and the other is interested in selling), the bank can offer these two clients the opportunity to trade internally without impacting the market price. Both clients, and the bank, benefit from this internalization. On the other hand, if the bank cannot find two matching clients, the bank has to perform the trade in the public market, which introduces some additional costs and might impact the price. Banks, therefore, put efforts into locating internalized matches.

One such effort is the following service. To incentivize clients to trade, banks publish a list of stocks that they are interested in trading, known as “*axe list*”. The *axe list* that the bank publishes contains, among other things, aggregated information on previous transactions that were made by clients and facilitated by the bank. For instance, to facilitate clients’ trades, the bank sometimes buys stocks that some clients wish to sell. The bank then looks to sell those stocks to other clients at advantageous rates before selling those stocks in the public market. Those stocks will appear in the bank’s *axe list*.

The *axe list* consists of tuples  $(op, symb, axe)$  where  $op \in \{buy, sell\}$ , *symb* is the symbol of the security to buy or sell, and *axe* is the number of the shares (quantity) of the security to buy or sell (we sometimes use the terminology of “*long*” for buy and “*short*” for sell). This *axe list* provides clients the ability to locate available (synthetic) trades at reduced financing rates.

Unfortunately, this method is unsatisfactory. Although the information in the *axe list* of the bank relates to transactions that were already executed, there is a correlation between previous transactions that a client performed and future transactions that it might wish to trade. Thus, clients feel uncomfortable with seeing their recent (potentially large) trade history (although anonymized and aggregated) in the *axe list* that

the bank publishes, and sometimes ask the bank to remove their previous trades from the axe list. Clients, therefore, face the following dilemma: keeping their axes published reveal information about their future potential trades or investment strategy, while continuously asking to remove trades from the axe list limits the banks’ ability to internalize trades and offer advantageous rates, to begin with.

The bank currently uses some ad-hoc methods to mitigate the leakage. For instance, it might aggregate several stocks together into “buckets” (e.g., reveal only the range of available stocks to trade in some sector), or trim the volumes of other stocks. This does not guarantee privacy, and also makes it harder to locate potential matches.

### 1.1 Our Work

We provide a novel method for addressing the inventory matching problem (a simple double auction, which is periodic, with a single fixed price per symbol). Our main contribution is a suite of cryptographic protocols for several variants of the inventory matching problem. The system we report, called Prime Match, was implemented and runs in production in J.P. Morgan since September 2021. Prime Match has the potential to transform common procedures in the financial world. We design the following systems:

- *Bank-to-client inventory matching:* Prime Match supports a secure two-party (bank-to-client) inventory matching. The client can privately find stocks to trade in the bank’s full axe list without the bank revealing its axe list, and without the client revealing the stocks and quantities it wishes to trade. The protocol is secure against a semi-honest bank and malicious client and is of two rounds of communication (three rounds if both parties learn the output).
- *Client-to-client inventory matching:* We extend Prime Match to support a secure (client-to-client) inventory matching. This is a three-party protocol where the bank is an intermediate party that mainly delivers the messages between two clients and facilitates the trade if there is a match. This enables two clients to explore whether they can have potential matches against each other and not just against the axe list of the bank. This further increases potential matches. The protocol is secure in the presence of one malicious corruption and is of three rounds of interaction.
- *Multiparty inventory matching:* We also extend the client-to-client inventory matching to multiple clients coming at once and looking to be matched.

We expand on each one of those scenarios below.

**Bank-to-client inventory matching:** We replace the current procedure in which the bank sends an axe list to a client, and the client replies with which stocks to trade based on the axe list, with a novel bank-to-client inventory matching. Prime Match allows the bank to locate potential matches without revealing its axe list, and without the client revealing its interests. Moreover, as the bank can freely use accurate

axe information (as the axe list is hidden), clients have no longer an interest to remove themselves from the axe list. All parties enjoy better internalization and advantageous rates.

Importantly, the bank does not learn any information about what the client is interested in on any stock that is not matched, and likewise, the client does not learn any information on what is available unless she/he is interested in that as well. Only after matches are found, the bank and the client are notified and the joint interest is revealed. At a high level, for two orders (buy,  $X$ ,  $axe_1$ ) and (sell,  $X$ ,  $axe_2$ ) on the same symbol  $X$ , we provide a secure two-party protocol that computes as the matching quantity the min quantity between  $axe_1$  and  $axe_2$ .<sup>1</sup>

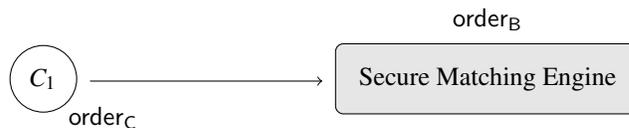


Figure 1: Client-to-bank topology. Client  $C$  sends an encrypted order  $order_C = (buy, X, axe_1)$  to the Bank (secure matching engine) which holds  $order_B = (sell, X, axe_2)$ . The engine computes the minimum between  $axe_1$  and  $axe_2$ .

**Client-to-client inventory matching:** The above approach only enables matching between the bank’s inventory to each client separately but does not allow a direct matching among different clients. For illustration, consider the following scenario: Client A is interested in buying 100 shares of some security  $X$ , while client B is interested in selling 200 shares of the same security  $X$ . On the other hand, the bank does not provide  $X$  in its inventory axe list. The bank either distributes in a non-private way its axe list to clients A and B (as it is being conducted prior to our work) or engages twice in a bank-to-client inventory matching described above, the first time against client A and the second time against client B. The two clients do not find  $X$  in the list, and both clients would have to trade on the public market at higher costs.

Prime Match allows the clients and the bank not to miss such opportunities. We provide a mechanism that acts as a transparent matching engine. Each client provides as input to the computation his/her encrypted axes, and the clients then interact and learn whether their axes match or not, see Figure 2. For this solution, we provide a three-party secure minimum protocol  $\Pi_{min}$  among two clients and the bank as the intermediary party to facilitate and execute the trade if there is a match.

**Multi-party protocol.** A potentially powerful mechanism would be to support multiple clients coming at the same time, where all clients talk to each other through the bank who

<sup>1</sup>In our actual protocol, each party also provides a range of quantities it wishes to trade, i.e., a minimum amount and a maximum amount. If there is no match that satisfies at least its minimum quantity, then there is no trade. To keep the introduction simple, we omit this additional complexity for now.



Figure 2: Client-to-client topology. Client  $C_1$  and  $C_2$  send encrypted orders  $\text{order}_1 = (\text{buy}, X, \text{axe}_1)$  and  $\text{order}_2 = (\text{sell}, X, \text{axe}_2)$ , respectively, to the Bank which computes the minimum of  $\text{axe}_1$  and  $\text{axe}_2$ .

facilitates the trades when there are matches. This might increase the potential number of matches for each client. We implement such a mechanism based on our client-to-client matching protocol, where we invoke it  $\binom{n}{2}$  times, for each possible pair of clients, when  $n$  is the total number of participating clients. Since the service of  $\text{axe}$  matching is relatively exclusive, i.e., it is offered only to selected clients,  $n$  is relatively small (around 10 per day), and thus this approach suffices for the current needs.

At this point, we only implement this relatively degenerated form of multiparty matching. We provide security for a semi-honest bank and malicious clients. In the multiparty setting, there are further challenges that have to be explored, such as what information is leaked by the functionality due to partial matches (i.e., client A can fulfill its order, say, selling 1000 shares by matching with client B and C, each wishing to buy 500 shares). Moreover, to achieve malicious security, the protocol also has to guarantee that the bank does not discriminate against clients, e.g., when two clients are both interested in buying some security, then it treats them fairly and does not prefer to match the “big client” over the “small client.” In fact, it is impossible to support security against a malicious bank in this case already because of the star network – the clients communicate through the bank with no authentication (see [8]). Therefore, achieving malicious security would require some different setups and further techniques. We leave this for future research.

From a business perspective, the clients generally do trust the bank, and the bank is also highly regulated and will not risk its reputation by attempting to cheat. Therefore, semi-honest bank generally suffices.

**Secure minimum protocol.** At the heart of our Prime Match engine is a secure protocol for comparing two input values  $\text{axe}_1$  and  $\text{axe}_2$ , each in  $\{0, \dots, 2^n - 1\} \subset \mathbb{F}_q$ . The protocol, given the bit-decompositions of  $\text{axe}_1$  and  $\text{axe}_2$ , computes the minimum between the two. We have a two-party variant (bank to client) and a three-party variant when only two parties have inputs (client-to-client) and the third party (the server) helps in the computation. For the latter, an interesting property of our protocol is that the two clients only perform linear operations, and therefore can operate non-interactively on encrypted inputs (or secret-shared, or homomorphic commitments, etc.). The server facilitates the computation. For  $\ell$ -bit inputs, our protocol runs in three rounds of interaction and with  $O(\ell^2)$  communication where in the first round clients provide their

input, and in the last (third) round the output is revealed. The protocols also offers malicious security.

**Implementation and evaluation.** All three scenarios were implemented, and we report running times in Section 5. On the bottom level, both bank-to-client and client-to-client protocols can process roughly 10 symbols per second with security against malicious clients under conventional machines with commodity hardware. Our system is running live, in production by J.P. Morgan. To the best of our knowledge, this is the first MPC solution running live in the financial world. Commercially, the main advantage of the system is the increased opportunities for clients to find matches.

As clients do not wish to spend resources to use such a service (installation of packages, maintenance cost, etc.), and cannot commit to providing tech resources before testing the product, Prime Match is implemented as a browser service. This raises several challenges in the implementation, see Section 5. Moreover, in the client-to-client matching a star topology network is required where clients communicate only with the bank. Clients do not wish to establish communication with other clients and reveal their identities to other clients.

**Our contributions.** To conclude, our contributions are:

- We identify a real-world problem in which cryptography significantly simplifies and improves the current inventory matching procedure.
- We provide two new protocols: bank-to-client inventory matching and client-to-client inventory matching. Those completely replace the current method which leaks information and misses potential matches. Our protocols are novel and are specifically tailored to the problem at hand. We do not just use generic, off-the-shelf, MPC protocols (see Section 1.3 for a discussion).
- At the heart of our matching engine is a novel two-round comparison protocol that minimizes interaction and requires only linear operations.
- The protocols are implemented and run live, in production, by a major bank in the US – J.P. Morgan.

## 1.2 Related Work

**Prior works on volume (quantity) matching.** We now compare the prior privacy-preserving volume matching architectures [5, 14, 18] to Prime Match. The MPC-based volume matching constructions of [14, 18] derive their security by separating the system’s *service operator/provider* into several (e.g., 3) distinct servers, whose collusion would void the system’s security guarantees. The clients submit their *encrypted* orders to the servers by secret sharing, such that no single server can recover the encrypted orders. The clients have no control over these servers and no clear way to prevent them from colluding.

Allowing clients to *themselves* serve as contributing operators of the system would present its own challenges. For instance, it would impose a disproportionate computational burden on those clients who choose to serve as operators. Moreover, it is unclear how to incentivize clients to run heavy computations, and to play the role of the operators.<sup>2</sup>

The fully homomorphic approach of [7] imposes a computational burden on a single server in a star topology network in which clients communicate with the server. Moreover, the concrete efficiency of the proposed GPU-FHE scheme is slow. Furthermore, the scheme of [7] does not offer malicious security. FHE-based solutions for malicious security are much less efficient than the ones based on MPC.

**Prior works on privacy-preserving dark pools.** A recent line of research has attempted to protect the information contained in dark pools [5, 15] run by an operator. The systems described in these works allow users to submit orders in an “encrypted” form; the markets’ operator then compares orders “through the encryptions”, unveiling them only if matches occur. The functionality of privacy-preserving dark pools is a continuous double auction in which apart from the *direction* (buy or sell) and a desired trading *volume*, a *price* (indicating the “worst” price at which the participant would accept an exchange) is submitted. The operator “matches” compatible orders, which, by definition, have opposite directions, and for which the price of the buy order (the “bid”) is at least the price of the sell order (the “ask”). [14, 15] are based on MPC with multiple operators and the work of [5] is based on FHE.

Dark pools are different than our setting, as matches are also conditioned on an agreement on a price (requiring many more comparisons) leading to more complex functionality. In comparison, inventory matching is a simple double auction, which is periodic, with a single fixed price per symbol. Moreover, dark pools support high-frequency trading, which means that they have to process orders very fast. All prior works’ performance on dark pools (including multi-server dark pools) does not suffice for high-frequency trading. In comparison, axe-list matching is a much slower process; with the current, insecure procedure of axe-matching, a few minutes might elapse between when the bank sends its axe list, and the time the client submits its orders. Since ensuring privacy introduces some overhead, clients might not necessarily prefer a slower privacy-preserving dark pool over a fast ordinary dark pool. Furthermore, secure comparison is a necessary building block for dark pools. Any of the comparison protocols from prior works, [16, 19, 28, 29, 32, 34], including ours, can be used for dark pools, but all of them have some overhead. Unfortunately, neither of these works can lead to a fast dark pool (in a star topology network) which is close to the running times of a dark pool operating on plaintexts. Achieving fast enough

<sup>2</sup>Part of the success of the Prime Match system is related to the fact that clients are offered a web service to participate in the system which requires minimal tech support by the clients.

comparison that is suitable for high-frequency trading is an interesting open problem.

The work of Massacci et al. [30] considers a distributed Market Exchange for futures assets which has functionality with multiple steps where one of the steps includes the dark pool functionality. Their experiments show that their system can handle up to 10 traders. Moreover, orders are not concealed: in particular, an aggregated list of all waiting buy and sell orders is revealed which is not the case in solution and the dark pool solutions. Note that there are works that propose dark pool constructions on the blockchain [6, 24, 31] which is not the focus of our work. Moreover, these solutions have different guarantees and security goals. None of the above solutions is in production.

### **Prior works on secure 3-party Less Than comparison.**

There are several works in the literature that propose secure comparison protocols of two values in the information-theoretic setting [2, 16, 19, 28, 29, 32, 34]. See Table 1 for a detailed comparison of these works compared to ours. Our protocol does not require preprocessing and runs in 2 rounds of interaction. Our cost incurs an  $\ell^2$  overhead since we secret share  $\ell$  bit numbers in a field of size  $\ell$ . Similar overhead also appears in prior works. The security parameter  $\lambda$  overhead is required due to the use of coin flipping and the additional use of commitments in the malicious protocol. The main reason for the higher overhead of prior secret sharing-based protocols in Table 1 is that they require interaction per secure multiplication leading to an increased round complexity ( $\approx \log \ell$ ). Our protocol does not require any secure multiplications, which is a significant benefit in upgrading our passive protocol to one with malicious security.

The works of [20, 21, 23, 27], based on multiplicative/additive homomorphic encryption, provide 2 (or constant) round solutions but they only offer passive security. The computational cost is capped at  $O(\lambda \cdot \ell)$  modular multiplications. Moreover, some works require a trusted setup assumption to generate the public parameters. For instance the modulus generation of the homomorphic Paillier encryption-based solutions.

The most recent work of [2], based on functional secret sharing in the preprocessing model, is a three-round solution offering only passive security with the cost of  $O(\ell)$  PRG calls in its online phase.

## **1.3 Why Specifically-Tailored Protocols?**

A natural question is why we design a specifically tailored protocol for the system, instead of just using any generic, off-the-shelf secure computation protocols. Those solutions are based on securely emulating arithmetic or Boolean circuits, and require translating the problem at hand to such a circuit. Specifically, for our client-to-client matching algorithm, which is a three-party secure protocol with one corruption, it looks promising to use some generic MPC protocols that are

Protocol	Offline Communication	Online Communication	Offline Computation	Online Computation	Rounds	Security	Corruption
[32]	-	$O((\ell \log \ell) \cdot (\ell + s))$	-	$O((\ell \log \ell) \cdot (\ell + s))$	31	passive	HM
[16]	-	$O(\ell \cdot (\ell + s))$	-	$O(\ell \cdot (\ell + s))$	$O(\log \ell)$	passive	HM
[34]	-	$O(\ell^2 + \log \ell)$	-	$O(\ell^2 + \log \ell)$	$O(\log \ell)$	passive	DM
This work	-	$O(\ell^2 + \log \ell)$	-	$O(\ell^2 + \log \ell)$	2	passive	DM
[28]	$O(\ell^2)$	$O(\log \ell \cdot (\ell + s))$	$O(\ell^2)$	$O(\log \ell \cdot (\ell + s))$	$O(\log \ell)$	active	HM
[19]	-	$O((\ell \log \ell) \cdot (\ell + s))$	-	$O((\ell \log \ell) \cdot (\ell + s))$	44	active	HM
[29]	$O(\ell)$	$O((\ell \log \ell) \cdot (\ell + s))$	$O(\ell)$	$O(\ell \cdot (\ell + s))$	$O(\log \ell)$	active	DM
This work	-	$O(\ell \cdot (\ell + \lambda))$	-	$O(\ell \cdot (\ell + \lambda))$	2	active	HM

Table 1: Cost of passive and active comparison protocols in terms of offline, and online communication and computation complexity; in terms of rounds; in terms of security; and in terms of corruptions supported. HM stands for honest majority, while DM stands for dishonest majority.  $\ell$  denotes to the bit length of the input,  $s$  is the statistical security parameter and  $\lambda$  is the computational security parameter. The work of [29] achieves statistical security over arithmetic fields but it achieves perfect security over the arithmetic rings.

based on replicated secret sharing, such as [3, 17] or garbled circuits [26, 35].

There are two main requirements from the system (from a business perspective) that leads us to design a specifically-tailored protocol and not a generic MPC: (1) The need for a constant number of rounds; (2) Working with committed inputs. Furthermore, no offline preprocessing is possible since clients wish to participate only during the live matching phase. We provide a comparison with generic MPC techniques in Appendix I.

## 1.4 Overview of our Techniques

n We focus in this overview on the task of client-to-client matching (see Figure 2): A three-party computation between two clients that communicate through the bank. We present our solution while hiding only the clients’ quantities axe. However, our detailed protocol additionally hides both the directions and the symbols. We present our protocol in the semi-honest setting and then explain how to achieve malicious corruption.

**Semi-honest clients and server:** The client provides secret shares (and commitments) for all possible symbols and for the two possible sides. If a client is not interested in buying (resp. selling) a particular stock, it provides 0 as its input for that symbol and size. It is assumed that the total number of symbols is around 1000 – 5000, and of course, the number of sides is 2. Thus, each party has to provide roughly 2000 – 10000 values. To see if there is a match between clients A and B on a particular stock, we securely compute the minimum between the values the parties provided with opposite sides (i.e., A sells and B buys, or B sells and A buys).

Each one of the clients first secret shares its secret value axe using an additive secret sharing scheme. The two clients then exchange shares<sup>3</sup>. Then, they decide on the matching

<sup>3</sup>The communication model does not allow the two clients to talk directly,

quantity by computing two bits indicating whether the two quantities are equal or which one of the two is minimal.

We design a novel algorithm for computing the minimum. The algorithm consists of two phases. As depicted in Figure 3, the first phase works on the shares of the two secrets (a, b), exchanged via the matching engine using symmetric key encryption, while performing only linear operations on them (min protocol). Looking ahead, each one of the two clients would run this phase, without any interaction, on its respective shares. The result would be shares of some secret state ( $d_0, d_1$ ) in which some additional non-linear processing is needed after reconstruction to obtain the final result. However, the secret state can be simulated with just the result of the computation – i.e., the two bits indicating whether the two numbers are equal or which one is minimal. Therefore, at the end of the first phase, the two clients can send the shares to the server, who reconstructs the secret state and learns the result, again using just local (this time, non-linear) computations.

Our minimum protocol min is described in Section 4.2, and we overview our techniques and contributions in the relevant section. Our semi-honest protocol is given in Section 4.3.

**Malicious clients.** We now discuss how to change the protocol to protect against malicious clients.

Zooming out from computing minimum, the auction works in two phases: a “registration” stage, where clients submit their orders, and the matching stage, where the clients and the bank run the secure protocols to find matching orders. In the malicious case, the parties submit commitments to the quantities of their orders to the server. The list of participants is not known in advance, only the clients who submitted a commitment can participate in the current matching phase. Moreover, the list of participants (at each run) is not public and is only known to the server.

and each client talks only to the server. However, using encryption and authentication schemes, the two clients can establish a secure channel while the server just delivers messages for them.

Of course, clients have to be consistent, and cannot use different values in the matching phase and in the registration phase. In the matching phase, the clients secret shares their inputs (additive secret sharing), and prove using a Zero-Knowledge (ZK) proof that the shares define the committed value provided in the registration phase.

More specifically, client  $C_1$  commits to  $a$  during registration, i.e., sends  $\text{Com}(a)$  to the server and commits to the shares  $(a_1, a_2)$  of the minimum by sending  $\text{Com}(a_1)$  and  $\text{Com}(a_2)$  to the server. It also proves in ZK the statement that  $\text{Com}(a) = \text{Com}(a_1) + \text{Com}(a_2)$  given that the commitment scheme is linearly homomorphic allowing to perform additions on committed values.

On top of the basic semi-honest protocol (as depicted in Figure 3) we also exchange the messages shown in Figure 4 where every party forwards a commitment to the other party for the share that it does not hold. Client  $C_1$  receives a commitment to  $b_2$  and client  $C_2$  receives a commitment to  $a_1$ .

Next, recall that our minimum protocol requires only linear work from the clients, and thus it allows to work on any linearly-homomorphic cryptosystem, such as linear secret sharing scheme, linearly homomorphic commitments, linearly homomorphic encryption scheme, and so on. In the semi-honest setting, we used this property to work only on the secret shares. We run the linear algorithm three times in parallel, on different inputs:

1. First, each client simply runs the algorithm on additive shares, just as in the semi-honest solution. This is depicted in Figure 3. Running the algorithm on those shares would result in shares of some secret state that will be delivered to the server. The server reconstructs the state and computes the result from this state.
2. Second, the parties run the algorithm on the commitments of the other party's share. This is depicted in Figure 4. Since the commitment scheme is also linearly-homomorphic, it enables Alice to compute a commitment of what Bob is supposed to send to the server in the first invocation, and vice versa.
3. Third, the parties also compute (again, using only linear operations!) information that allows the server to learn the openings of the other party's commitment. This enables the server to check that all values it received in the first invocations are correct.

**Malicious server.** Our final system does not provide security against a malicious server, unless the two clients can authenticate themselves to each other, or can talk directly. We show that if the client can communicate to each other, then we can also support malicious server for the comparison protocol.

The server receives shares of some secret states, together with commitments of the secret states. It then reconstructs the secret state and checks for consistency. It then has to perform some non-linear operations on the secret state to

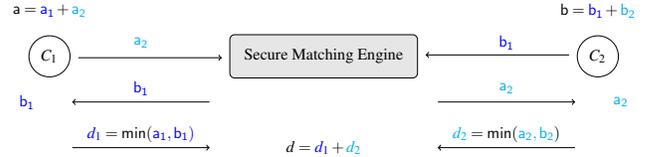


Figure 3: Client-to-client matching protocol for computing the minimum between the quantities  $a$  from client  $C_1$  and  $b$  from client  $C_2$  in the semi-honest setting. As described in Footnote 3, the communication between the two clients through the server is encrypted, and so the view of the server in this communication is just  $d_1, d_2$ .

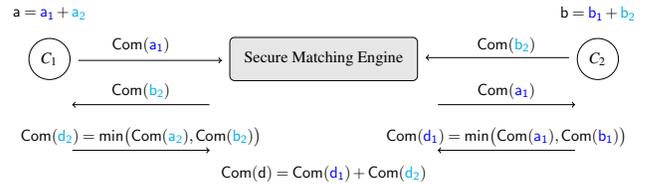


Figure 4: Client-to-client matching protocol for computing the minimum in the presence of a malicious adversary. In addition to values computed in Figure 3, the parties compute commitments of the value that the other participant is supposed to send to the server.

learn the result. Applying generic ZK proofs for proving that the non-linear operation was done correctly would increase the overhead of our solution. Luckily, the non-linear operation that the server performs is ZK-friendly, specifically, it is enough to prove in ZK a one-out-of-many proof (i.e., given a vector, proving that one of the elements in the vector is zero; see Theorem 4.4). For this particular language, there exists fast ZK solutions [25]. See Section 4.4 for a description and details.

**Organization.** The paper is organized as follows. In Section 2 we provide the preliminaries, while some are deferred to the appendices. In Section 3 we provide the main matching engine functionality. In Section 4 we provide our protocol for computing the minimum, including the semi-honest and the malicious versions. In Section 5 we report the system performance and in Appendix J we mention challenges pertaining to the deployment of our system.

## 2 Preliminaries

Some preliminaries are deferred to Appendix A.

**Notations.** We use PPT as an acronym for probabilistic polynomial time. We use  $\lambda$  to denote the security parameter, and  $\text{negl}(\lambda)$  to denote a negligible function (a function that is smaller than any polynomial for sufficiently large  $\lambda$ ).

**Commitment schemes.** A commitment scheme is a pair of probabilistic algorithms  $(\text{Gen}, \text{Com})$ ; given public key

parameters  $\text{params} \leftarrow \text{Gen}(1^\lambda)$  and a message  $m$ ,  $\text{com} := \text{Com}(\text{params}, m; r)$  returns a “commitment” to the message  $m$ . To reveal  $m$  as an opening of  $\text{com}$ , the committer simply sends  $m$  and  $r$  (this is sometimes called “decommitment”). For notational convenience, we often omit  $\text{params}$ . A commitment scheme is *homomorphic* if, for each  $\text{params}$ , its message, randomness, and commitment spaces are abelian groups, and the corresponding commitment function is a group homomorphism. We always write message and randomness spaces additively and write commitment spaces multiplicatively. See Section A for more details.

**Zero-knowledge proofs.** We use non-interactive zero-knowledge for three languages. See formal treatment in Appendix A.1:

- **Commitment equality proof:** Denoted as the relation  $\mathcal{R}_{\text{ComEq}}$ , the prover convinces the verifier that the two given commitments  $c_0, c_1$  hide the same value.
- **Bit proof:** Denoted as the relation  $\mathcal{R}_{\text{BitProof}}$ , allows the prover to prove that a commitment  $c$  hides a bit, i.e., a value in  $\{0, 1\}$ .
- **Out-of-many proofs.** Denoted as the relation  $\mathcal{R}_{\text{OneMany}}$ , allows a prover to prove that one of the commitments  $V_0, \dots, V_n$  in the statement is a commitment of 0.

### 3 The Prime Match Main Functionalities

We now describe our Prime Match inventory matching functionalities. We describe the bank-to-client functionality (Section 3.1), the client-to-client functionality (Section 3.2), and the multi-client system (Section 3.3).

#### 3.1 Bank to Client Matching

This variant is a two-party computation between a bank and a client. The bank tries to find matching orders between its own inventory and each client separately. As mentioned in the introduction, this essentially comes to replace the current procedure of axe-matching as being conducted today, with a privacy-preserving mechanism. Today the bank sends its inventory list to the client who then submits orders to the bank. Note, however, that if the bank runs twice with two different clients, and the bank does not hold some security  $x$ , and two clients are interested in  $x$  with opposite directions, then such a potential match will not be found.

The functionality proceeds as follows. The client sends to the bank its axe list. This includes the list of securities it is interested in, and for each security whether it is interested in long (buy) or short (sell) exposures, and the quantity. The client sends its own list. The functionality finds whether the bank and the client are interested in the same securities with opposite sides, and in that case, it provides as output the matching orders and the quantity is the minimum between the two amounts.

#### FUNCTIONALITY 3.1 ( $\mathcal{F}_{\text{B2C}}$ –Bank-to-client functionality).

The functionality is parameterized by the set of all possible securities to trade, a set  $U$ .

**Input:** The bank  $P^*$  inputs lists of orders  $(\text{symb}_i^*, \text{side}_i^*, \text{amount}_i^*)$  where  $\text{symb}_i \subseteq U$  is the security,  $\text{side}_i^* \in \{\text{buy}, \text{sell}\}$  and  $\text{amount}_i^*$  is an integer. The client sends its list of the same format,  $(\text{symb}_i^C, \text{side}_i^C, \text{amount}_i^C)$ .

**Output:** Initialize a list of Matches. For each  $i, j$  such that  $\text{symb}_i^* = \text{symb}_j^C$  and  $\text{side}_i^* \neq \text{side}_j^C$ , add  $(\text{symb}_i^*, \text{side}_i^*, \text{side}_j^C, \min\{\text{amount}_i^*, \text{amount}_j^C\})$  to  $M$ . Output  $M$  to both parties.

From a business perspective, it is important to note that the input of the client (and the bank) serves as a “commitment” - if a match is found then it is executed right away.

The functionality resembles a set intersection. In set intersection, if some element is in the input set of some party but not in its output, it can conclude that it does not contain in the other set. Here, if a party does not find a particular symbol in its output although it did provide it as an input, then it is still uncertain whether the other party is not interested in that security, or whether it is interested but with the same side. We show how to implement the functionality in the presence of a malicious client or a semi-honest server in Appendix E.

**Bank to multiple clients.** In the actual system, the bank has to serve multiple clients. This is implemented by a simple (sequential) composition of the functionality. Specifically, the functionality is now reactive where clients first register that they are interested to participate. The bank then runs Functionality 3.1 with the clients – either according to the basis of first-come-first-served, or some random ordering. We omit the details and exact formalism as they are quite natural given a semi-honest bank.

**Range functionalities.** In Appendix G, we show a variant of the protocol where each party inputs a *range* in which it is interested and not just one single value. I.e., if a matched order does not satisfy some minimal value, it will not be executed. Since the minimum value does not change throughout the execution, whenever the bank receives 0 as a result of the execution it cannot decide whether the client is not interested in that particular symbol, or whether it is interested – but the matched amount does not satisfy the minimum threshold.

#### 3.2 Client to Client Matching

In this variant, the bank has no input and it tries to find potential matches between clients facilitating two clients that wish to compare their inventories. This is a three-party computation where the bank just facilitates the interaction. It is important to notice that the clients do not know each other,

and do not know who they are being paired with. The bank selects the two clients and offers them to be paired.

**FUNCTIONALITY 3.2** ( $\mathcal{F}_{C2C}$ –Client-to-client functionality).

The functionality is parameterized by the set of all possible securities to trade, a set  $U$ . This is a three-party functionality between two clients,  $P_1$  and  $P_2$ , and the bank  $P^*$ .

**Input:** The client  $P_1$  inputs a list of orders  $(\text{symb}_i^1, \text{side}_i^1, \text{amount}_i^1)$ . The client  $P_2$  inputs a list of orders  $(\text{symb}_i^2, \text{side}_i^2, \text{amount}_i^2)$ , and the bank has no input.

**Output:** Initialize a list of Matches. For each  $i, j$  such that  $\text{symb}_i^1 = \text{symb}_j^2$  and  $\text{side}_i^1 \neq \text{side}_j^2$ , add  $(\text{symb}_i^1, \text{side}_i^1, \text{side}_j^2, \min\{\text{amount}_i^1, \text{amount}_j^2\})$  to  $M$ . Output  $M$  to all three parties.

In the next section, we show how to implement this functionality in the presence of a malicious client or a malicious server, assuming that the two clients can communicate directly, or have a public-key infrastructure. When the two clients can communicate only through the server and there is no public-key infrastructure (PKI) or any other setup, there is no authentication and the server can impersonate the other client. We therefore cannot hope to achieve malicious security. We achieve security in the presence of a semi-honest server. See a discussion in the next subsection.

### 3.3 The Multi-Client System

We now proceed to the multiparty auction. Here we have parties that register with their intended lists, and the bank facilitates the orders by pairing the clients according to some random order. The functionality is now reactive; The parties first register, in which they announce that they are willing to participate in the next auction, and they also commit to their orders. In the second phase, the bank selects pairs of clients in a random order to perform client-to-client matching. Looking ahead, typically there are around 10 clients that participate in a given auction.

For simplicity of exposition and to ease readability, we write the functionality as the universe is just a single symbol. Moreover, instead of sending the side explicitly, the client sends two integers  $L$  and  $S$ , representing its interest in long (buy) or short (sell) exposure, respectively. Rationally, each party would put one of the integers as 0 (as otherwise, it would just pay extra fees to the bank). Generalizing the functionality to deal with many symbols is performed in a natural manner, where the number of total symbols is 1000-5000 in practice. The main functionality can process all the different symbols in parallel.

**FUNCTIONALITY 3.3** ( $\mathcal{F}_{MC}$  – Multi-client matching). This is an  $n + 1$  party functionality between  $n$  clients  $P_1, \dots, P_n$  and a bank  $P^*$ .

Upon initialization,  $\mathcal{F}_{MC}$  initializes a list  $\mathcal{P} = \emptyset$  and two vectors  $\mathcal{L}$  and  $\mathcal{S}$  of size  $n$ , where  $n$  bounds the total number of possible clients.

$\mathcal{F}_{MC}$ .**Register**( $P_i, L_i, S_i$ ). Store  $\mathcal{L}[i] = L_i$  and  $\mathcal{S}[i] = S_i$  and add  $i$  to  $\mathcal{P}$ . Send to the bank  $P^*$  the message registered( $P_i$ ).

$\mathcal{F}_{MC}$ .**Process**().

1. Choose a random ordering  $O$  over all pairs of  $\mathcal{P}$ .
2. For the next pair  $(i, j) \in O$  try to match between  $P_i$  and  $P_j$  (we can assume wlog that always  $i \leq j$ ):
  - (a) Compute  $M_0 = \min(\mathcal{L}[i], \mathcal{S}[j])$ ,  $b_0^0 = (\mathcal{L}[i] \leq \mathcal{S}[j])$ ,  $b_1^0 = (\mathcal{S}[i] \leq \mathcal{L}[j])$ .
  - (b) Compute  $M_1 = \min(\mathcal{S}[i], \mathcal{L}[j])$ ,  $b_0^1 = (\mathcal{S}[i] \leq \mathcal{L}[j])$  and  $b_1^1 = (\mathcal{S}[i] \leq \mathcal{L}[j])$ .
  - (c) Send  $(i, j, M_0, M_1, b_i^0, b_i^1)$  to  $P_i$ , and  $(i, j, M_0, M_1)$  with  $(b_0^0, b_1^0, b_0^1, b_1^1)$  to  $P^*$ .
  - (d) Update  $\mathcal{L}[i] = \mathcal{L}[i] - M_0$  and  $\mathcal{S}[j] = \mathcal{S}[j] - M_0$ .
  - (e) Update  $\mathcal{S}[i] = \mathcal{S}[i] - M_1$  and  $\mathcal{L}[j] = \mathcal{L}[j] - M_1$ .

**On malicious server.** Our final protocol (see Appendix F) for  $\mathcal{F}_{MC}$  is secure in the presence of a semi-honest  $P^*$  (and a malicious client). Inherently, clients communicate through a star network where the bank facilitates the communication. Moreover, we assume no PKI, clients do not know how many clients are registered in the system, and how many clients are participating in the current auction. This can be viewed as “secure computation without authentication”, in which case the server can always “split” the communication and disconnect several parties from others (see [8] for a formal treatment).

We prove security in the presence of a semi-honest server. In fact, our protocol achieves a stronger notion of a guarantee than just semi-honest, as in particular, it runs the underlying comparison protocol (a single invocation of a client-to-client matching) which is secure against a malicious server.

Another relaxation that we make is that the ordering of pairs is random, and we do not have a mechanism to enforce it. Note also that the functionality leaks some information to the server; in particular, after finding a match, the bank executes it immediately. The bank can infer information about whether two values equal to 0, and therefore whether a client is not interested in a particular symbol. In contrast, each client just learns whether its value is smaller or equal to the value of the other party, and therefore when it inputs 0 it can never infer whether the other party is interested in that symbol or not.

## 4 Securely Computing Minimum

A pivotal building block in Prime Match is a secure minimum protocol. In Section 4.1, we review our functionality for computing the minimum. We focus on the case of client-to-client matching with an aiding server. We show how to convert the protocol for two parties in Appendix E.

In Section 4.2 we present the underlying idea for computing the minimum. The algorithm computes the minimum while using only linear operations (looking ahead, those would be computed on shared values) while pushing the non-linear operations on reconstructed data. In Sections 4.3 and 4.4 we show a semi-honest and a malicious protocol for computing the minimum, respectively.

### 4.1 The Minimum Functionality

After receiving a secret integer from each one of the two parties, the functionality compares them and gives as a result two bits – which indicate which one of the two inputs is smaller than the other, or whether they are equal. It also gives the result to the server.

**FUNCTIONALITY 4.1** ( $\mathcal{F}_{\text{comp}}$ : Server-aided secure minimum functionality).

Consider two players,  $P_0$  and  $P_1$ , and a server  $P^*$ .

- **Input:**  $P_0$  and  $P_1$  respectively send integers  $v_0$  and  $v_1$  in  $\{0, \dots, 2^n - 1\}$  to  $\mathcal{F}_{\text{comp}}$ .
- **Output:**  $\mathcal{F}_{\text{comp}}$  sends  $b_0 := (v_0 \leq v_1)$  to  $P_0$ ,  $b_1 := (v_1 \leq v_0)$  to  $P_1$ ,  $(b_0, b_1)$  to  $P^*$ .

In the rest of this section, we will show how to implement this functionality in the presence of a semi-honest (Section 4.3) and malicious adversary (Section 4.4).<sup>4</sup>

### 4.2 Affine-Linear Comparison Function

We first describe an abstract algorithm which compares two elements  $v_0$  and  $v_1$  of  $\{0, \dots, 2^n - 1\} \subset \mathbb{F}_q$ , given their bit-decompositions. We separate the algorithm into two parts: ComparisonInitial (Algorithm 1) and ComparisonFinal (Algorithm 2). Both parts do not use any underlying cryptographic primitives.

In the first algorithm (ComparisonInitial), all operations on the bit-decompositions of the two inputs  $v_0$  and  $v_1$  are *linear*. Looking ahead, this will be extremely useful when converting the algorithm into a secure two-party protocol, where  $v_0$  and  $v_1$  are additively shared between the two parties (or also just committed, encrypted under additively homomorphic encryption scheme, etc.). In particular, this part of the protocol can be executed without any interaction, just as the algorithm

<sup>4</sup>For ease of presentation, Functionality 4.1 is for the semi-honest version of the protocol; For the malicious case, we will use a slightly different functionality on committed inputs; See Appendix D.

itself when  $v_0$  and  $v_1$  are given in the clear. The second algorithm (ComparisonFinal) can be computed by a *different* party, given all information in the clear. Looking ahead, this will be executed by the server  $P^*$  on the outputs of the first part. This part contains some non-linear operations, however, this part of the algorithm does not have to be translated into a secure protocol.

**Overview Algorithm 1** (ComparisonInitial). Our approach is inspired by the algorithm of Wagh, Gupta, and Chandran [34, Alg. 3], which compares a *secret-shared* integer with a *public* integer. (Specifically, its inputs consist of an array of public bits and an array of secret-shared bits.) We extend the algorithm and allow the comparison of two private integers using only linear operations.

We achieve this  $\mathbb{F}_q$ -linearity in the following way. We fix integers  $v_0$  and  $v_1$  in  $\{0, \dots, 2^n - 1\}$ , with big-endian bit-decompositions given respectively by

$$v_0 = \sum_{j=0}^{n-1} 2^{n-1-j} \cdot v_{0,j}, \quad \text{and} \quad v_1 = \sum_{j=0}^{n-1} 2^{n-1-j} \cdot v_{1,j}.$$

We follow the paradigm of [34], whereby, for each  $j \in \{0, \dots, n-1\}$ , a quantity  $w_j$  is computed which equals 0 if and only if  $v_{0,j} = v_{1,j}$ . Meanwhile, for each  $j \in \{0, \dots, n-1\}$ , we set  $c_j := 1 + v_{0,j} - v_{1,j} + \sum_{k < j} w_k$  (we also set  $c_n := \sum_{k=0}^{n-1} w_k$ , as we discuss below). The crucial observation of [34] is that, for each  $j \in \{0, \dots, n-1\}$ ,  $c_j = 0$  so long as  $v_{0,j} < v_{1,j}$  as bits (that is, if  $1 + v_{0,j} - v_{1,j} = 0$ ) *and* the higher bits of  $v_0$  and  $v_1$  agree (inducing the equality  $\sum_{k < j} w_k = 0$ ). By consequence, *some*  $c_j$ , for  $j \in \{0, \dots, n-1\}$ , must equal 0 whenever  $v_0 < v_1$ . Similarly,  $c_n = 0$  whenever  $v_0 = v_1$ . In summary,  $v_0 \leq v_1$  implies that some  $c_j = 0$ , for  $j \in \{0, \dots, n\}$ .

The main challenge presented by this technique is to ensure that the opposite implication holds; that is, we must prevent the sum  $c_j := 1 + v_{0,j} - v_{1,j} + \sum_{k < j} w_k$  from equalling 0 (possibly by overflowing) modulo  $q$ —that is, even when  $w_k \neq 0$  for some  $k < j$ —and hence yielding a “false positive”  $c_j = 0$ , which would falsely assert the inequality  $v_0 \leq v_1$ . [34] prevents this phenomenon by ensuring that each  $w_j \in \{0, 1\}$ , and choosing  $2 + n < q$  (they set  $n = 64$  and  $q = 67$ ). In fact, [34] defines  $w_j := (v_{0,j} - v_{1,j})^2$ . Under this paradigm,  $c_j := 1 + v_{0,j} - v_{1,j} + \sum_{k < j} w_k$  is necessarily *non-zero* so long as either  $v_{0,j} \geq v_{1,j}$  as bits (so that  $1 + v_{0,j} - v_{1,j} > 0$ ) or *any*  $w_k \neq 0$ , for  $k < j$ .

This squaring operation is nonlinear in the bits  $v_{0,j}$  and  $v_{1,j}$ , and so it is unsuitable for our setting. We adopt the following recourse instead, which yields  $\mathbb{F}_q$ -linearity at the cost of requiring that the number of bits  $n \in O(\log q)$  (a mild restriction in practice). The key technique is that we may eliminate the squaring—thereby allowing each  $w_j$  to remain in  $\{-1, 0, 1\}$ —provided that we multiply each  $w_j$  by a suitable public scalar. In fact, it suffices to multiply each (unsquared) difference  $w_j$  by  $2^{2+j}$ . In Theorem 4.4 below, we argue that this approach is correct.

Our modifications to [34] also include our computation of the *non-strict* inequality  $v_0 \leq v_1$ —effected by the extra value  $c_n$ —as well as our computation of the opposite non-strict inequality,  $v_1 \leq v_0$ , in parallel. The latter computation proceeds identically, except uses  $-1 + v_{0,j} - v_{1,j} \in \{-2, -1, 0\}$  at each bit.

---

**Algorithm 1** ComparisonInitial  $((v_{0,0}, \dots, v_{0,n-1}), (v_{1,0}, \dots, v_{1,n-1}))$

---

- 1: Assign  $w_{\text{accum}} := 0$
  - 2: **for**  $j \in \{0, \dots, n-1\}$  **do**
  - 3:     Set  $c_{0,j} := 1 + v_{0,j} - v_{1,j} + w_{\text{accum}}$ .
  - 4:     Set  $c_{1,j} := -1 + v_{0,j} - v_{1,j} + w_{\text{accum}}$
  - 5:     Set  $w_j := (v_{0,j} - v_{1,j})$  and  $w_{\text{accum}} += 2^{2+j} \cdot w_j$
  - 6: Set  $c_{0,n}$  and  $c_{1,n}$  equal to  $w_{\text{accum}}$
  - 7: Sample a random permutation  $\pi \leftarrow \mathbf{S}_{n+1}$
  - 8: **for**  $j \in \{0, \dots, n\}$  **do**
  - 9:     Sample random scalars  $s_{0,j}, s_{1,j}$  from  $\mathbb{F}_q \setminus \{0\}$ .
  - 10:     Assign  $d_{0,j} := s_{0,j} \cdot c_{0,\pi(j)}$ ,
  - 11:     Assign  $d_{1,j} := s_{1,j} \cdot c_{1,\pi(j)}$
  - 12: **return**  $(d_{0,0}, \dots, d_{0,n})$  and  $(d_{1,0}, \dots, d_{1,n})$
- 

Of course, the intermediate value  $v_{0,j} - v_{1,j}$  need only be computed once per iteration of the first loop.

**Overview of Algorithm 2** (ComparisonFinal). Note that in Algorithm 1,  $w_{\text{accum}} = 0$  as long as  $v_{0,j} = v_{1,j}$ , and it attains a non-zero value at the first  $j$  for which  $v_{0,j} \neq v_{1,j}$ . Up to that point,  $(c_{0,j}, c_{1,j}) = (1, -1)$ . At the first  $j$  for which  $v_{0,j} \neq v_{1,j}$ :

- If  $v_0 > v_1$  (i.e.,  $(v_{0,j}, v_{1,j}) = (1, 0)$ ), then we get that  $(c_{0,j}, c_{1,j}) = (2, 0)$ .
- If  $v_0 < v_1$  (i.e.,  $(v_{0,j}, v_{1,j}) = (0, 1)$ ), then we get that  $(c_{0,j}, c_{1,j}) = (0, -2)$ .

The algorithm then makes sure that no other value of  $c_{b,j'} = 0$ , essentially by assigning  $w_{\text{accum}}$  to be non-zero. If  $v_0 = v_1$  then  $c_{0,n} = c_{1,n} = 0$ . Finally, all the bits  $(c_{0,0}, \dots, c_{0,n}), (c_{1,0}, \dots, c_{1,n})$  are permuted and re-randomized with some random scalars. Observe that if  $v_0 > v_1$  then all the values  $d_{0,0}, \dots, d_{0,n}$  are all non-zero, and one of  $d_{1,0}, \dots, d_{1,n}$  is zero. If  $v_1 \geq v_0$  then exactly one of the values  $d_{1,0}, \dots, d_{1,n}$  is 0 and all  $d_{0,0}, \dots, d_{0,n}$  are non-zero.

It is crucial that the vectors  $(d_{0,0}, \dots, d_{0,n})$  and  $(d_{1,0}, \dots, d_{1,n})$  do not contain any information on  $v_0, v_1$  rather than whether  $v_0 \leq v_1$  or  $v_1 \leq v_0$ . Specifically, these values can easily be simulated given just the two bits  $v_0 \leq v_1$  and  $v_1 \leq v_0$ . Therefore, it is safe to give both vectors to a third party, which will perform the non-linear part of the algorithm. For a vector of bits  $(x_0, \dots, x_n) \in \{0, 1\}^{n+1}$ , the operation  $\mathbf{any}_{j=0}^n x_j$  returns 1 iff there exists  $j \in [0, \dots, n]$  such that  $x_j = 1$ . Algorithm 2 simply looks for the 0 coordinate in the two vectors. We have:

---

**Algorithm 2** ComparisonFinal  $((d_{0,0}, \dots, d_{0,n}), (d_{1,0}, \dots, d_{1,n}))$

---

- 1: Assign  $b_0 := \mathbf{any}_{j=0}^n (d_{0,j} = 0)$
  - 2: Assign  $b_1 := \mathbf{any}_{j=0}^n (d_{1,j} = 0)$
  - 3: **return**  $b_0$  and  $b_1$
- 

In the below theorem, we again consider bit-decomposed integers  $v_0 = \sum_{j=0}^{n-1} 2^{n-1-j} \cdot v_{0,j}$  and  $v_1 = \sum_{j=0}^{n-1} 2^{n-1-j} \cdot v_{1,j}$ ; we view the bits  $v_{i,j}$  as elements of  $\{0, 1\} \subset \mathbb{F}_q$ . The following theorem is proven in Appendix B:

**Theorem 4.4.** *Suppose  $n$  is such that  $2 + 4 \cdot (2^n - 1) < q$ . Then for every  $v_0, v_1 \in \mathbb{F}_q$*

$$(v_0 \leq v_1, v_1 \leq v_0) = \text{ComparisonFinal}(\text{ComparisonInitial}(\vec{v}_0, \vec{v}_1)) ,$$

where  $\vec{v}_0, \vec{v}_1$  are the bit-decomposition of  $v_0, v_1$ , respectively. Moreover, for every  $i \in \{0, 1\}$ :

- If  $v_i \leq v_{1-i}$  then there exists exactly one  $j \in \{0, \dots, n\}$  such that  $d_{i,j} = 0$ . Moreover,  $j$  is distributed uniformly in  $\{0, \dots, n\}$ , and each  $d_{i,k}$  for  $k \neq j$  is distributed uniformly in  $\mathbb{F}_q \setminus \{0\}$ .
- If  $v_0 > v_{1-i}$  then the vector  $(d_{i,0}, \dots, d_{i,n})$  is distributed uniformly in  $\mathbb{F}_q^{n+1} \setminus \{0\}^{n+1}$ .

We emphasize that Algorithm 1 uses only  $\mathbb{F}_q$ -linear operations throughout. A number of our below protocols conduct Algorithm 1 “homomorphically”; that is, they execute the algorithm on elements of an  $\mathbb{F}_q$ -module  $M$  which is unequal to  $\mathbb{F}_q$  itself. As a basic example, Algorithm 1 may be executed on bits  $(v_{0,j})_{j=0}^{n-1}$  and  $(v_{1,j})_{j=0}^{n-1}$  which are *committed*, provided that the commitment scheme is homomorphic (its message, randomness and commitment spaces should be  $\mathbb{F}_q$ -modules, and its commitment function an  $\mathbb{F}_q$ -module homomorphism). Furthermore, Algorithm 1 may be conducted on additive  $\mathbb{F}_q$ -shares of the bits  $(v_{0,j})_{j=0}^{n-1}$  and  $(v_{1,j})_{j=0}^{n-1}$ .

In this latter setting, sense must be given to the affine additive constants  $\pm 1$ . As in [34], we specify that these be shared in the obvious way; that is, we stipulate that  $P_0$  and  $P_1$  use the shares 0 and  $\pm 1$ , respectively.

### 4.3 The Semi-Honest Protocol

For simplicity, we first describe a protocol that securely computes this functionality in the setting of three-party computation with an honest majority and a *semi-honest* adversary. We give a maliciously secure version in Protocol 4.3 and prove its security in Section C of the Appendix.

**Theorem 4.5.** *If  $\Pi_{\text{CT}}$  is a secure coin-tossing protocol,  $G$  is a pseudorandom generator, and the two clients communicate using symmetric authenticated encryption with pseudorandom*

**PROTOCOL 4.2** (Semi-honest secure comparison protocol).

- **Input:**  $P_0$  and  $P_1$  hold integers  $v_0$  and  $v_1$ , respectively, in  $\{0, \dots, 2^n - 1\}$ .  $P^*$  has no input.
- **The protocol:**
  1.  $P_0$  and  $P_1$  engage in the coin-tossing procedure  $\Pi_{CT}$  (see Sect. A) to obtain a  $\lambda$ -bit shared secret  $s$ .
  2. Each party  $P_i$  (for  $i \in \{0, 1\}$ ) computes the bit decomposition  $v_i = \sum_{j=0}^{n-1} 2^{n-1-j} \cdot v_{i,j}$ , for bits  $v_{i,j} \in \{0, 1\}$ .
  3. For each  $j \in \{0, \dots, n-1\}$ ,  $P_i$  computes a random additive secret-sharing  $v_{i,j} = \langle v_{i,j} \rangle_0^q + \langle v_{i,j} \rangle_1^q$  in  $\mathbb{F}_q$ .  $P_i$  sends the shares  $(\langle v_{i,j} \rangle_{1-i}^q)_{j=0}^{n-1}$  to  $P_{1-i}$ .
  4. After receiving the shares  $(\langle v_{1-i,j} \rangle_i^q)_{j=0}^{n-1}$ , from  $P_{1-i}$ ,  $P_i$  executes Algorithm 1 on the appropriate shares; that is, it evaluates

$$\left( (\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n \right) \leftarrow \text{ComparisonInitial} \left( (\langle v_{0,j} \rangle_i^q)_{j=0}^{n-1}, (\langle v_{1,j} \rangle_i^q)_{j=0}^{n-1} \right),$$

where all internal random coins are obtained from  $G(s)$ .

5.  $P_i$  sends the output shares  $(\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n$  to  $P^*$ .
6. After receiving all shares,  $P^*$  reconstructs for every  $j \in \{0, \dots, n\}$ :

$$d_{0,j} := \langle d_{0,j} \rangle_0^q + \langle d_{0,j} \rangle_1^q \quad \text{and} \quad d_{1,j} := \langle d_{1,j} \rangle_0^q + \langle d_{1,j} \rangle_1^q,$$

and finally executes Algorithm 2 to receive  $(b_0, b_1) := \text{ComparisonFinal} \left( (d_{0,j})_{j=0}^n, (d_{1,j})_{j=0}^n \right)$ .

7.  $P^*$  sends  $b_i$  to  $P_i$ .

- **Output:** Each  $P_i$  outputs  $b_i$ .  $P^*$  outputs  $(b_0, b_1)$ .

ciphertexts, then Protocol 4.2 securely computes Functionality 4.1 in the presence of a semi-honest adversary corrupting at most one party. Each party sends or receives  $O(n^2 + \lambda)$  bits, where  $n$  is the length of the input and  $\lambda$  is the security parameter.

## 4.4 The Maliciously Secure Protocol

We now give our malicious protocol for Functionality 4.1 in Protocol 4.3. To ease notation, we denote  $N = \{0, \dots, n-1\}$ . We already gave an overview of the protocol as part of the introduction. The parties commit to the inputs, share their inputs, commit to the shares, and prove that all of those are consistent. Then, each party can operate on the shares it received (as in the semi-honest protocol), but also on the commitment that it holds on the other parties' share. When  $P^*$  receives the shares of  $d$  from party  $P_i$ , it also receives a commitment of what  $P_{1-i}$  is supposed to send. The server can therefore check for consistency, and that no party cheated. Moreover, each party  $P_i$  can also compute commitments of the two vectors  $\vec{d}_0, \vec{d}_1$ . When the server comes to prove  $P_i$  that  $v_i \leq v_{1-i}$ , it has to show that there is a 0 coordinate in its vector  $\vec{d}_i$ . This is possible using one-out-of-many proof. See Appendix D for the full security proof. We note that the functionality is slightly different than that of Functionality 4.1.

## 4.5 Bank-to-Client

We show in Appendix E a two-party (Bank-to-Client) version of the protocol, where the bank does not just facilitate the computation but also provides input. In a nutshell, the parties use linear homomorphic encryption – ElGamal encryption – instead of secret sharing.

## 5 Prime Match System Performance

We report benchmarks of Prime Match in two different environments, a Proof of Concept (POC) environment, and the production environment after refactoring the code to meet the requirements of the bank's systems. The former benchmarks can be used to value the performance of the comparison protocol for other applications in different systems.

**Secure Minimum Protocol Performance:** For the purposes of practical convenience, adoption, and portability, our client module is entirely browser-based, written in JavaScript. Its cryptographically intensive components are written in the C language with side-channel resistance, compiled using Emscripten into WebAssembly (which also runs natively in the browser). Our server is written in Python, and also executes its cryptographically intensive code in C. Both components are multi-threaded—using WebWorkers on the client side and a thread pool on the server's—and can execute arbitrarily many concurrent instances of the protocol in parallel (i.e., constrained only by hardware). All players communicate by

**PROTOCOL 4.3** (Maliciously secure comparison protocol  $\Pi_{\text{comp}}$ ).**Input:**  $P_0$  and  $P_1$  hold integers  $v_0$  and  $v_1$ , respectively, in  $\{0, \dots, 2^n - 1\}$ .  $P^*$  has no input.**Setup phase:** A coin-tossing protocol  $\Pi_{\text{CT}}$ , and a commitment scheme (Gen, Com), are chosen (see Sect. 2,A).**The protocol:**

1. Commit  $V_i \leftarrow \text{Com}(v_i; r_i)$ , and send  $V_i$  to  $P^*$ .  $P^*$  delivers  $V_i$  to  $P_{1-i}$ .
2. Engage with  $P_{1-i}$  in the coin-tossing procedure  $\Pi_{\text{CT}}$ , and obtain a  $\lambda$ -bit shared  $s$ .
3. Compute the bit decomposition  $v_i = \sum_{j \in N} 2^{n-1-j} \cdot v_{i,j}$ , for bits  $v_{i,j} \in \{0, 1\}$ .
4. For each  $j \in N$ :
  - (a) Compute a random additive secret-sharing  $v_{i,j} = \langle v_{i,j} \rangle_0^q + \langle v_{i,j} \rangle_1^q$  in  $\mathbb{F}_q$ .
  - (b) Commit  $V_{i,j,k} \leftarrow \text{Com}(\langle v_{i,j} \rangle_k^q; r_{i,j,k})$  for each  $k \in \{0, 1\}$ .
  - (c) Open  $V_{i,j,1-i}$  by sending  $\langle v_{i,j} \rangle_{1-i}^q$  and  $r_{i,j,1-i}$  to  $P_{1-i}$ .
5. Send the full array  $(V_{i,j,k})_{j,k=0}^{n-1,1}$  to  $P_{1-i}$ .
6. Compute:  $\pi_i \leftarrow \text{ComEq.Prove}(V_i, \prod_{j=0}^{n-1} (V_{i,j,0} \cdot V_{i,j,1})^{2^{n-1-j}})$ ,  $\pi_{i,j} \leftarrow \text{BitProof.Prove}(V_{i,j,0} \cdot V_{i,j,1})$ , for all  $j \in N$ ,  $P_i$  sends  $\pi_i$ , and  $(\pi_{i,j})_{j=0}^{n-1}$  to  $P_{1-i}$ .
7.  $P_i$  receives  $\pi_{1-i}$ , and  $(\pi_{1-i,j})_{j=0}^{n-1}$  and verifies the following:
  - (a) The openings  $\langle v_{1-i,j} \rangle_i^q$  and  $r_{1-i,j,i}$  indeed open  $V_{1-i,j,i}$ , for  $j \in N$ ,
  - (b)  $\text{ComEq.Verify}(\pi_{1-i}, V_{1-i}, \prod_{j=0}^{n-1} (V_{1-i,j,0} \cdot V_{1-i,j,1})^{2^{n-1-j}})$ ,
  - (c)  $\text{BitProof.Verify}(\pi_{1-i,j}, V_{1-i,j,0} \cdot V_{1-i,j,1})$  for each  $j \in N$ .
 If any of these checks fail,  $P_i$  aborts.
8.  $P_i$  runs Algorithm 1, in parallel, on the shares  $\langle v_{k,j} \rangle_i^q$ , the randomnesses  $r_{k,j,i}$ , and the commitments to the *other party's* shares  $V_{k,j,1-i}$  (all for  $j \in N$  and  $k \in \{0, 1\}$ ). That is,  $P_i$  runs:

$$\begin{aligned} \left( (\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n \right) &\leftarrow \text{ComparisonInitial} \left( (\langle v_{0,j} \rangle_i^q)_{j=0}^{n-1}, (\langle v_{1,j} \rangle_i^q)_{j=0}^{n-1} \right), \\ \left( (s_{0,j,i})_{j=0}^n, (s_{1,j,i})_{j=0}^n \right) &\leftarrow \text{ComparisonInitial} \left( (r_{0,j,i})_{j=0}^{n-1}, (r_{1,j,i})_{j=0}^{n-1} \right), \\ \left( (D_{0,j,1-i})_{j=0}^n, (D_{1,j,1-i})_{j=0}^n \right) &\leftarrow \text{ComparisonInitial} \left( (V_{0,j,1-i})_{j=0}^{n-1}, (V_{1,j,1-i})_{j=0}^{n-1} \right), \end{aligned}$$

using the same shared randomness  $s$  for all internal coin flips.

9.  $P_i$  sends  $(\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n$  and the randomnesses  $(s_{0,j,i})_{j=0}^n, (s_{1,j,i})_{j=0}^n$  to  $P^*$ .

**Party  $P^*$  (Output reconstruction):** After receiving all shares,  $P^*$  proceeds as follows:

1. Reconstruct for each  $j \in N$ :

$$\begin{aligned} d_{0,j} &:= \langle d_{0,j} \rangle_0^q + \langle d_{0,j} \rangle_1^q, & s_{0,j} &:= s_{0,j,0} + s_{0,j,1} \\ d_{1,j} &:= \langle d_{1,j} \rangle_0^q + \langle d_{1,j} \rangle_1^q, & s_{1,j} &:= s_{1,j,0} + s_{1,j,1} \end{aligned}$$

2. Finally,  $P^*$  executes Algorithm 2, that is:  $(b_0, b_1) := \text{ComparisonFinal} \left( (d_{0,j})_{j=0}^n, (d_{1,j})_{j=0}^n \right)$
3. For each  $i \in \{0, 1\}$ , if  $b_i$  is true, then  $P^*$  re-commits  $D_{i,j} := \text{Com}(d_{i,j}; s_{i,j})$  for each  $j \in N$ , computes  $\pi'_i \leftarrow \text{OneMany.Prove} \left( (D_{i,j})_{j=0}^n \right)$ , and finally sends  $\pi'_i$  to  $P_i$ . Otherwise,  $P^*$  sends  $\perp$  to  $P_i$ .  $P^*$  outputs both  $b_0$  and  $b_1$ .

**Each Party  $P_i$  (output reconstruction):**

1. If receives a proof  $\pi'_i$ , compute  $D_{i,j} := \text{Com}(\langle d_{i,j} \rangle_i^q; s_{i,j,i}) \cdot D_{i,j,1-i}$ , for each  $j \in N$ , and then verifies  $\text{OneMany.Verify} \left( \pi'_i, (D_{i,j})_{j=0}^n \right)$ . If verification passes then output true, otherwise, output false.

	Number of Symbols	Latency (sec)	Throughput (transactions/sec)	Matching received msg Size (MB)	Matching sent msg Size (MB)
Bank-to-client	100	9.903 ( $\pm 0.174$ )	10.09	0.215	0.521
	200	19.533 ( $\pm 0.530$ )	10.23	0.430	1.040
	500	46.223 ( $\pm 0.787$ )	10.81	1.076	2.597
	1000	95.396 ( $\pm 2.063$ )	10.48	2.152	5.194
	2000	183.186 ( $\pm 2.512$ )	10.91	4.304	10.382
	4000	356.740 ( $\pm 2.149$ )	11.21	8.608	20.762
	10000	941.813 ( $\pm 18.465$ )	10.61	21.520	51.902
Client-to-client	100	11.15 ( $\pm 0.060$ )	8.96	0.972	1.549
	200	20.636 ( $\pm 0.525$ )	9.69	1.945	3.096
	500	51.493 ( $\pm 2.343$ )	9.71	4.863	7.737
	1000	101.051 ( $\pm 2.587$ )	9.89	9.727	15.472
	2000	208.813 ( $\pm 1.479$ )	9.57	19.454	30.942
	4000	390.510 ( $\pm 3.020$ )	10.24	38.908	61.882
	10000	1064.443 ( $\pm 70.439$ )	9.39	97.270	154.702

Table 2: Performance of Bank-to-Client matching for two clients, and Client-to-Client matching.

sending binary data on WebSockets (all commitments, proofs, and messages are serialized).

We run our experiments on commodity hardware throughout since our implementation is targeted to a real-world application where clients hold conventional computers. In particular, one of the two clients runs on an Intel Core i7 processor, with 6 cores, each 2.6GHz, and another one runs on an Intel Core i5, with 4 cores, each 2.00 GHz. Both of them are Windows machines. Our server runs in a Linux AWS instance of type c5a.8xlarge, with 32 vCPUs. In the first scenario, we run the client-to-bank inventory matching protocol for two clients where we process each client one by one against the bank’s inventory. In Table 2 we report the performance for a different number of registered orders per client (100, 200, . . . , 10000). *Latency* refers to the total time it takes to process all the orders from both clients (in seconds). *Throughput* measures the number of transactions per second. The number of orders/symbols processed per second is approximately 10. We also report the message size (in MB) for the message sent from each client to the server during the registration phase and the matching phase. We further record the size of the messages received from the server to each client. The bandwidth is 300 Mbps. Note that the message sizes can be reduced considerably if message serialization is not required.

In the second scenario, we run the client-to-client inventory matching protocol (the comparison protocol) for two clients where we try to match the orders of the two clients via the bank. In this case we assume that the bank has no inventory. We report the performance in Table 2. The number of orders/symbols processed per second is approximately 10 in this case too.

**Prime Match Performance in Production:** Figure 6 shows a sketch of both the bank’s network tiering and the application architecture of Prime Match. The right side of the diagram shows that the clients are able to access the Prime Match UI

through the bank’s Markets portal with the correct entitlement.

The application UI for Prime Match is hosted on the bank’s internal cloud platform. After each client is getting authenticated by the bank’s Markets portal, it can access the application on the browser which establishes web socket connections to the server. The client’s computation takes place locally within their browser, ensuring that all private data remains local. The server is hosted on the bank’s trade management platform.

The bank’s network tiering exists to designate the network topology in and between approved security gateways (firewalls). The network traffic from the client application is handled by tier 1 which helps the bank’s Internet customers achieve low latency, secured, and accelerated content access to internally hosted applications. Then, the traffic will be subsequently forwarded to a web socket tunnel in tier 2 and gets further directed to the server in tier 3.

During the axe registration phase, the client logs in to Prime Match from his/her web browser and uploads a file of orders (symbol, directions, quantity) which are encrypted locally and securely on the browser. The client is uploading a minimum (threshold) and a maximum (full) quantity per symbol. Prime Match first processes the encrypted threshold quantities of all clients and then it processes the full quantities. This process implements the range functionality of Appendix G. Note that for the partial matches, the full quantity is never revealed to the server. Moreover, if there is no match no quantity is revealed to the server. Figure 7 shows a complete run of an auction on the client’s browser after the matching phase where the final matched quantities of some test symbols/securities with the corresponding fixed spread are revealed.

Our protocol runs in production every 30 minutes. There are two match runs each hour and matching starts at xx:10 and xx:40. Axe registration starts 8 minutes before matching (xx:02 and xx:32). The matching process finishes at various

times (as shown in the experiments section) based on the number of symbols. Based on the application requirements (up to 5000 symbols in the US and 10-60 clients) it does not finish after xx:55.

For running the code in production in the bank’s environment, the two clients run on the same type of Windows machines whose specs are Intel XEON CPU E3-1585L, with 4 cores, each 3.00 GHz. The server runs on the bank’s trade management platform with 32 vCPUs. In Table 3 of the Appendix, we report the performance. Moreover, we discuss challenges we faced during the implementation in Appendix J.

## 6 Conclusion

Inventory matching is a fundamental service in the traditional financial world. In this work, we introduce secure multiparty computation in financial services by presenting a solution for matching orders in a stock exchange while maintaining the privacy of the orders. Information is revealed only if there is a match. Our central tool is a new protocol for secure comparison with linear operations in the presence of a malicious adversary, which can be of independent interest. Our system is running live, in production, and is adopted by a large bank in the US – J.P. Morgan.

## 7 Acknowledgements

We would like to thank Mike Reich, Vaibhav Papat, Sitaraman Rajamani, Dan Stora, James Mcilveen, Oluwatoyin Aguiyi, Niall Campbell, Wanyi Jiang, Grant McKenzie, Steven Price, Vinay Gayakwad, Srikanth Veluvolu, Noel Peters for their great efforts and help to move Prime Match in production. Last but not least, we would like to thank our executive sponsor Jason Sippel. The paper is based on joint patents [4, 22]

This paper was prepared in part for information purposes by the Artificial Intelligence Research group and AlgoCRYPT CoE of JPMorgan Chase & Co and its affiliates (“JP Morgan”), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful. 2023 JP Morgan Chase & Co. All rights reserved.

## References

- [1] ‘Bristol Fashion’ MPC Circuits. <https://homes.esat.kuleuven.be/~nsmart/MPC/>, accessed: December 2022.
- [2] Amit Agarwal, Stanislav Peceny, Mariana Raykova, Phillipp Schoppmann, and Karn Seth. Communication efficient secure logistic regression. *Cryptology ePrint Archive*, Paper 2022/866, 2022. <https://eprint.iacr.org/2022/866>.
- [3] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 805–817. ACM, 2016.
- [4] Gilad Asharov, Tucker Balch, Hans Buehler, Richard Hua, Antigoni Polychroniadou, and Manuela Veloso. Systems and methods for privacy-preserving inventory matching, filed Dec. 6, 2019.
- [5] Gilad Asharov, Tucker Hybinette Balch, Antigoni Polychroniadou, and Manuela Veloso. PPDPs: Privacy-preserving dark pools. In *19th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2020)*, 2020. Extended abstract.
- [6] Samiran Bag, Feng Hao, Siamak F Shahandashti, and Indranil Ghosh Ray. Seal: Sealed-bid auction without auctioneers. *IEEE Transactions on Information Forensics and Security*, 15:2042–2052, 2019.
- [7] Tucker Balch, Benjamin E. Diamond, and Antigoni Polychroniadou. Secretmatch: Inventory matching from fully homomorphic encryption. In *Proceedings of the First ACM International Conference on AI in Finance, ICAIF ’20*, New York, NY, USA, 2020. Association for Computing Machinery.
- [8] Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. *J. Cryptol.*, 24(4):720–760, 2011.
- [9] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas P. Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael I. Schwartzbach, and Tomas Toft. Secure multiparty computation goes live. In *Financial Cryptography and Data Security, FC 2009*, volume 5628, pages 325–343. Springer, 2009.
- [10] Benedikt Bünz, Shashank Agrawal, Mahdi Zamani, and Dan Boneh. Zether: Towards privacy in a smart contract world. In *International Conference on Financial Cryptography and Data Security*, 2020. Full version.

- [11] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy (SP)*, volume 1, 2018. Full version.
- [12] Matteo Campanelli, Rosario Gennaro, Steven Goldfeder, and Luca Nizzardo. Zero-knowledge contingent payments revisited: Attacks and payments for services. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017*, pages 229–243. ACM, 2017.
- [13] Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
- [14] John Cartlidge, Nigel P. Smart, and Younes Talibi Alaoui. Mpc joins the dark side. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*, pages 148–159. Association for Computing Machinery, 2019. Full version.
- [15] John Cartlidge, Nigel P. Smart, and Younes Talibi Alaoui. Multi-party computation mechanism for anonymous equity block trading: A secure implementation of turquoise plato uncross. Cryptology ePrint Archive, Report 2020/662, 2020. <https://ia.cr/2020/662>.
- [16] Octavian Catrina and Sebastiaan de Hoogh. Improved primitives for secure multiparty integer computation. In *SCN*, 2010.
- [17] Koji Chida, Daniel Genkin, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Yehuda Lindell, and Ariel Nof. Fast large-scale honest-majority MPC for malicious adversaries. In *Advances in Cryptology - CRYPTO 2018*, volume 10993, pages 34–64. Springer, 2018.
- [18] Mariana Botelho da Gama, John Cartlidge, Antigoni Polychroniadou, Nigel P. Smart, and Younes Talibi Alaoui. Kicking-the-bucket: Fast privacy-preserving trading using buckets. *Financial Cryptography*, 2022.
- [19] Ivan Damgård, Matthias Fitzi, Eike Kiltz, Jesper Buus Nielsen, and Tomas Toft. Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In *Theory of Cryptography, TCC 2006*, volume 3876, pages 285–304. Springer, 2006.
- [20] Ivan Damgård, Martin Geisler, and Mikkel Krøigaard. Homomorphic encryption and secure comparison. *Int. J. Appl. Cryptogr.*, 1(1):22–31, 2008.
- [21] Ivan Damgård, Martin Geisler, and Mikkel Krøigaard. A correction to ‘efficient and secure comparison for online auctions’. *Int. J. Appl. Cryptogr.*, 1(4):323–324, 2009.
- [22] Benjamin Diamond and Antigoni Polychroniadou. Privacy-preserving inventory matching with security against malicious adversaries, filed Jan. 28, 2021.
- [23] Marc Fischlin. A cost-effective pay-per-multiplication comparison method for millionaires. In *CT-RSA 2001*, volume 2020, pages 457–472. Springer, 2001.
- [24] Hisham S Galal and Amr M Youssef. Publicly verifiable and secrecy preserving periodic auctions. In *International Conference on Financial Cryptography and Data Security*, pages 348–363. Springer, 2021.
- [25] Jens Groth and Markulf Kohlweiss. One-out-of-many proofs: Or how to leak a secret and spend a coin. In *Advances in Cryptology – EUROCRYPT 2015*, volume 9057 of *Lecture Notes in Computer Science*, pages 253–280. Springer Berlin Heidelberg, 2015.
- [26] Jonathan Katz, Samuel Ranellucci, Mike Rosulek, and Xiao Wang. Optimizing authenticated garbling for faster secure two-party computation. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 365–391. Springer, 2018.
- [27] Hsiao-Ying Lin and Wen-Guey Tzeng. An efficient solution to the millionaires’ problem based on homomorphic encryption. In *Applied Cryptography and Network Security, Third International Conference, ACNS 2005, New York, NY, USA, June 7-10, 2005, Proceedings*, volume 3531 of *Lecture Notes in Computer Science*, pages 456–466, 2005.
- [28] Helger Lipmaa and Tomas Toft. Secure equality and greater-than tests with sublinear online complexity. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part II*, volume 7966 of *Lecture Notes in Computer Science*, pages 645–656. Springer, 2013.
- [29] Eleftheria Makri, Dragos Rotaru, Frederik Vercauteren, and Sameer Wagh. Rabbit: Efficient comparison for secure multi-party computation. In *Financial Cryptography and Data Security - FC 2021*, volume 12674, pages 249–270. Springer, 2021.
- [30] Fabio Massacci, Chan Nam Ngo, Jing Nie, Daniele Venturi, and Julian Williams. Futuresmex: Secure, distributed futures market exchange. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 335–353. IEEE Computer Society, 2018.

- [31] Chan Nam Ngo, Fabio Massacci, Florian Kerschbaum, and Julian Williams. Practical witness-key-agreement for blockchain-based dark pools financial trading.
- [32] Takashi Nishide and Kazuo Ohta. Constant-round multiparty computation for interval test, equality test, and comparison. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, 90-A(5):960–968, 2007.
- [33] Björn Terelius and Douglas Wikström. Efficiency limitations of  $\Sigma$ -protocols for group homomorphisms revisited. In *Security and Cryptography for Networks*, pages 461–476, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [34] Sameer Wagh, Divya Gupta, and Nishanth Chandran. SecureNN: 3-party secure computation for neural network training. In *Proceedings on Privacy Enhancing Technologies*, volume 2019, pages 26–49, 2019.
- [35] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 21–37. ACM, 2017.

## A Deferred Preliminaries

**Notations.** A distribution ensemble  $X = \{X(a, \lambda)\}_{a \in \mathcal{D}, \lambda \in \mathbb{N}}$  is an infinite sequence of random variables indexed by  $a \in \mathcal{D}$  and  $\lambda \in \mathbb{N}$ . We let  $x \leftarrow X(a, \lambda)$  denote sampling an element  $x$  according to the distribution  $X(a, \lambda)$ . Two distribution ensembles  $X = \{X(a, \lambda)\}_{a \in \mathcal{D}, \lambda \in \mathbb{N}}$  and  $Y = \{Y(a, \lambda)\}_{a \in \mathcal{D}, \lambda \in \mathbb{N}}$  are computationally-indistinguishable, denoted  $X \approx_{\text{comp}} Y$ , if for every non-uniform PPT algorithm  $D$  there exists a negligible function  $\text{negl}(\cdot)$  such that for every  $\lambda$  and  $a \in \mathcal{D}$ :

$$\Pr[D(X(a, \lambda)) = 1] - \Pr[D(Y(a, \lambda)) = 1] \leq \text{negl}(\lambda).$$

We let  $\mathcal{G}$  denote a *group-generation algorithm*, which on input  $1^\lambda$  outputs a cyclic group  $\mathbb{G}$ , its prime order  $q$  (with bit-length  $\lambda$ ) and a generator  $g \in \mathbb{G}$ . We recall the discrete-logarithm assumption, in which given a cyclic group  $(\mathbb{G}, q, g)$  and a random  $h \in \mathbb{G}$ , any probabilistic polynomial time adversary finds  $x$  such that  $g^x = h$  with at most negligible probability in  $\lambda$ . The Decisional Diffie-Hellman assumption states that any probabilistic polynomial time adversary distinguishes between  $(\mathbb{G}, q, g, g^x, g^y, g^{xy})$  and  $(\mathbb{G}, q, g, g^x, g^y, g^z)$  for random  $x, y, z \in \mathbb{F}_q$  with at most negligible probability in  $\lambda$ .

**Secure multiparty computation.** We use standard definitions of secure multiparty computation, following the stand-alone model [13]. We distinguish between security in the presence of a semi-honest adversary, i.e., the adversary follows the

protocol execution but looks at the all messages it received and tries to learn some additional information about the honest parties’ inputs, and a malicious adversary, i.e., an adversary that might act arbitrarily. We compare between a function  $\mathcal{F}(x_1, \dots, x_n) = (y_1, \dots, y_n)$  and a protocol  $\pi(x_1, \dots, x_n)$  that allegedly privately computes the function  $\mathcal{F}$ .

We review some notations: Let  $\mathcal{A}$  be a non-uniform probabilistic polynomial-time adversary controlling parties  $I \subset [n]$ . Let  $\text{REAL}_{\pi, \mathcal{A}(z), I}(x_1, \dots, x_n, \lambda)$  denote the output of the honest parties and  $\mathcal{A}$  in an real execution of  $\pi$ , with inputs  $x_1, \dots, x_n$ , auxiliary input  $z$  for  $\mathcal{A}$ , and security parameter  $\lambda$ . Let  $\mathcal{S}$  be a non-uniform probabilistic polynomial-time adversary. Let  $\text{IDEAL}_{\mathcal{F}, \mathcal{S}(z), I}(x_1, \dots, x_n, \lambda)$  denote the output of the honest parties and  $\mathcal{S}$  in an ideal execution with the functionality  $\mathcal{F}$ , inputs  $x_1, \dots, x_n$  to the parties, auxiliary-input  $z$  to  $\mathcal{S}$ , and security parameter  $\lambda$ . The security definition requires that for every real-world adversary  $\mathcal{A}$ , there exists an ideal world adversary  $\mathcal{S}$  such that  $\text{REAL}$  and  $\text{IDEAL}$  are indistinguishable. We refer to [13] for the full definition.

**Hybrid model and composition.** The hybrid model combines both the real and ideal models. Specifically, an execution of a protocol  $\pi_{\mathcal{F}}$  in the  $\mathcal{G}$ -hybrid model, for some functionality  $\mathcal{G}$ , involves the parties sending normal messages to each other (as in the real model), but they also have access to a trusted party computing  $\mathcal{G}$ . The composition theorem of [13] implies that if  $\pi_{\mathcal{F}}$  securely computes some functionality  $\mathcal{F}$  in the  $\mathcal{G}$ -hybrid model and a protocol  $\pi_{\mathcal{G}}$  computing  $\mathcal{G}$  then the protocol  $\pi_{\mathcal{F}}^{\mathcal{G}}$ , where every ideal call of  $\mathcal{G}$  is replaced with an execution of  $\pi_{\mathcal{G}}$  securely-computes  $\mathcal{F}$  in the real world.

**On modeling security.** We sometimes wish to prove that the protocol is secure with, e.g., malicious  $P_0$  but semi-honest  $P_1$ . We use the malicious security formalization where we quantify for adversaries that control  $P_1$  that do follow the protocol specifications. Note that in the ideal world, the simulator receives the input of the corrupted parties, whereas, in the malicious setting, we usually extract the input it actually uses in the protocol. In the case of a semi-honest adversary, the simulator does not have to extract the input, and it can just send the input it received to the trusted party.

**Coin-tossing.** We define the coin-tossing functionality,  $\mathcal{F}_{\text{CT}}$ . The functionality assumes two parties,  $P_1$  and  $P_2$ :

1. **Input:** Each party inputs  $1^\lambda$  and some parameter  $\ell > \text{sec}$ .
2. **The functionality:** Sample a uniform string  $s \leftarrow \{0, 1\}^\ell$  and give the two parties  $s$ .

This functionality can be realized easily, assuming commitment schemes and pseudorandom generators (security with abort).

**Commitment Schemes.** The hiding property of the commitment scheme states that for any two messages  $m_0, m_1$ , any adversary distinguishes between  $\text{Com}(\text{params}, m_0; r_0)$

and  $\text{Com}(\text{params}, m_1; r_1)$  with at most negligible probability in  $\lambda$ . The binding property states that given  $\text{params} \leftarrow \text{Gen}(1^\lambda)$ , any adversary can find  $(m_0, r_0), (m_1, r_1)$  such that  $\text{Com}(\text{params}, m_0; r_0) = \text{Com}(\text{params}, m_1; r_1)$  with at most negligible probability in  $\lambda$ .

For example, we have the Pedersen commitment scheme (see, e.g., [11, Def. 6]). For  $\text{params}$  consisting of a prime  $q$ , a cyclic group  $\mathbb{G}$  of order  $q$ , and elements  $g$  and  $h$  of  $\mathbb{G}$ , the Pedersen scheme maps elements  $m$  and  $r$  of  $\mathbb{F}_q$  to  $\text{Com}(m; r) := g^m h^r$ . If the discrete logarithm problem is hard with respect to  $\mathbb{G}$ , then the Pedersen scheme is computationally binding.

## A.1 Zero Knowledge Proofs

We present definitions for zero-knowledge arguments of knowledge, closely following [25] and [11]. We also review a number of *particular* constructions.

We posit a triple of interactive, probabilistic polynomial time algorithms  $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ . We fix a polynomial-time-decidable ternary relation  $\mathcal{R} \subset (\{0, 1\}^*)^3$ ; by definition, each common reference string  $\sigma \leftarrow \text{Setup}(1^\lambda)$  yields an NP language  $L_\sigma = \{x \mid \exists w : (\sigma, x, w) \in \mathcal{R}\}$ . We denote by  $\text{tr} \leftarrow \langle \text{Prove}(\sigma, x, w), \text{Verify}(\sigma, x) \rangle$  the (random) transcript of an interaction between  $\text{Prove}$  and  $\text{Verify}$  on auxiliary inputs  $(\sigma, x, w)$  and  $(\sigma, x)$  (respectively). Abusing notation, we occasionally write  $b \leftarrow \langle \text{Prove}(\sigma, x, w), \text{Verify}(\sigma, x) \rangle$  for the single bit indicating the verifier’s state upon completing the interaction (i.e., *accept* or *reject*). The zero-knowledge proof constructions we use (see below) all feature relations  $\mathcal{R}$  whose *statements* consist of commitments and whose *witnesses* consist of openings to these commitments. To simplify notation—we adopt a somewhat non-standard notational scheme, whereby we omit the witness from each call to  $\text{Prove}$ .

**The Zero-Knowledge Functionality  $\mathcal{F}_{\text{ZK}}(\mathcal{R})$ :** The zero-knowledge functionality  $\mathcal{F}_{\text{ZK}}(\mathcal{R})$  for an NP relation  $\mathcal{R}$  is defined as follows:

- **Input:** The prover and the verifier hold the same setup string  $\sigma$  and input  $x$ , and the prover holds  $w$ .
- **Output:** If  $\mathcal{R}(\sigma, x, w) = 1$  then it sends *accept* to the verifier. Otherwise, it sends *reject*.

The *particular* zero-knowledge proof constructions we use all feature relations  $\mathcal{R}$  whose *statements* consist of commitments and whose *witnesses* consist of openings to these commitments. More precisely—for each  $\mathcal{R}$  we consider below, and for each setup string  $\sigma$ —each element  $x$  of  $L_\sigma$  consists of one more commitment, while any valid witness  $w$  to  $x$ ’s membership consists exactly of openings (i.e., message and randomness) to some or all among these commitments.

In light of this fact—and to simplify notation—we adopt a somewhat non-standard notational scheme, whereby we omit the witness from each call to  $\text{Prove}$ , leaving it implicit in the

statement. For example, for a language  $L_\sigma$  whose statements consist of single commitments, we write  $\Pi.\text{Prove}(V)$ —for some commitment  $V$ —to mean  $\Pi.\text{Prove}(\sigma, V, (m, r))$ , where  $V = \text{Com}(m; r)$  (and the commitment scheme  $(\text{Gen}, \text{Com})$  is implicit in the setup string  $\sigma$ ). The expression  $\Pi.\text{Verify}(\pi, V)$  of course retains its usual meaning. Informally, we view each commitment  $V$  as a “data structure” whose internal “fields”  $m$  and  $r$  may, optionally (that is, for the prover), be populated with an opening to  $V$ . The call  $\Pi.\text{Prove}(V)$  acts on  $V$  by accessing its internal fields.

Moreover, given commitments  $V_0$  and  $V_1$  whose internal openings are populated, we understand  $V_0 \cdot V_1$  as a third commitment whose openings are *also* populated (in the obvious way, i.e., by addition).

**Commitment equality proof.** A simple Schnorr proof can be used to show that two Pedersen commitments open to the same message (recall that in Pedersen’s commitment, the commitment scheme is  $\text{Com}(m; r) = g^m h^r$  for two group elements  $g, h$ ). We specialize  $(\text{Gen}, \text{Com})$  to the Pedersen scheme, and set  $\sigma$  to a Pedersen base  $(g, h)$ . For completeness, we outline the details. The relevant relation is as follows:

$$\begin{aligned} \mathcal{R}_{\text{ComEq}} = \{ & (\sigma, V_0, V_1) \mid \exists (m_0, r_0, m_1, r_1) \\ & \text{s.t. } V_i = \text{Com}(m_i; r_i) \forall i \in \{0, 1\} \wedge m_0 = m_1 \}. \end{aligned}$$

The relation  $\mathcal{R}_{\text{ComEq}}$  expresses knowledge of openings to two commitments whose messages are *equal*. The protocol is trivial; it essentially runs a Schnorr protocol on the difference between  $V_0$  and  $V_1$ . This protocol can be derived as a special case of the “basic protocol” of Terelius and Wikström [33, Prot. 1]; it also appears implicitly in the techniques of [10, Sec. G]. The prover sends  $K = h^k$  for a random  $k \leftarrow \mathbb{F}_q$ . The verifier chooses a random  $x \leftarrow \mathbb{F}_q$ . The prover replies with  $s := (r_0 - r_1) \cdot x + k$  and the verifier accepts iff  $h^s = (V_0 \cdot V_1^{-1})^s \cdot K$ .

**Bit proof.** We now recall two important protocols from Groth, and Kohlweiss [25]. The first, [25, Fig. 1], demonstrates knowledge of a secret opening  $(m, r)$  of a public commitment  $V$  for which  $m \in \{0, 1\}$ . We use the protocol  $\text{BitProof}$  from [25]. The bit-proof protocol concerns the following relation:

$$\begin{aligned} \mathcal{R}_{\text{BitProof}} = \{ & (\sigma, V) \mid \exists (m, r), \text{ s.t. } V = \text{Com}(m; r) \\ & \wedge m \in \{0, 1\} \}. \end{aligned}$$

**One-out-of-many proofs.** The second protocol—namely, [25, Fig. 2]—demonstrates knowledge of a secret *element*  $V_l$  of a public list  $(V_0, \dots, V_n)$ —as well as a secret randomness  $r$ —for which  $V_l = \text{Com}(0; r)$ . Though [25] uses the former protocol merely as a subroutine of the latter, we will have occasion to make *independent* use of both protocols. For a fixed integer

$n$ , we have the relation:

$$\mathcal{R}_{\text{OneMany}} = \{(\sigma, (V_0, \dots, V_n)) \mid \exists l \in [0, \dots, n], r, \\ \text{s.t. } V_l = \text{Com}(0; r)\},$$

as well as the proof protocol OneMany, which we use exactly as written in [25, Fig. 2] without modification.

In practice, we choose  $n$  which is one short of a power of 2 (i.e.,  $n = 2^m - 1$  for some  $m$ ). The reason for this unusual indexing scheme will become clear below.

## B Proof of Theorem 4.4

**Theorem B.1** (Theorem 4.4, restated). *Suppose  $n$  is such that  $2 + 4 \cdot (2^n - 1) < q$ . Then for every  $v_0, v_1 \in \mathbb{F}_q$*

$$(v_0 \leq v_1, v_1 \leq v_0) = \\ \text{ComparisonFinal}(\text{ComparisonInitial}(\vec{v}_0, \vec{v}_1)),$$

where  $\vec{v}_0, \vec{v}_1$  are the bit-decomposition of  $v_0, v_1$ , respectively. Moreover, for every  $i \in \{0, 1\}$ :

- If  $v_i \leq v_{1-i}$  then there exists exactly one  $j \in \{0, \dots, n\}$  such that  $d_{i,j} = 0$ . Moreover,  $j$  is distributed uniformly in  $\{0, \dots, n\}$ , and each  $d_{i,k}$  for  $k \neq j$  is distributed uniformly in  $\mathbb{F}_q \setminus \{0\}$ .
- If  $v_0 > v_{1-i}$  then the vector  $(d_{i,0}, \dots, d_{i,n})$  is distributed uniformly in  $\mathbb{F}_q^{n+1} \setminus \{0\}^{n+1}$ .

*Proof.* We recall that Algorithms 1 and 2 are correct so long as, for each  $j \in \{0, \dots, n\}$ , the quantities  $c_{0,j}$  and  $c_{1,j}$  equal 0 if and only if, respectively,  $v_{0,j} < v_{1,j}$  and  $v_{0,j} > v_{1,j}$  (if  $j < n$ ) and moreover it holds that  $v_{0,k} = v_{1,k}$  for each  $k < j$ .

To show the nontrivial “only if” implications, it suffices to show that, for each  $j$ , the quantities  $c_{0,j}$  and  $c_{1,j}$  are nonzero so long as either some  $w_k \neq 0$ , for  $k < j$ , or else the appropriate bit inequality (i.e., either  $v_{0,j} < v_{1,j}$  or  $v_{0,j} > v_{1,j}$ ) fails to hold. The special case in which all  $w_k$  are 0 is trivially settled; in this case,  $c_{0,j}$  and  $c_{1,j}$  are easily seen to be nonzero when  $v_{0,j} \geq v_{1,j}$  and  $v_{0,j} \leq v_{1,j}$ , respectively (this is guaranteed so long as  $-2$  and  $2$  are unequal to 0 in  $\mathbb{F}_q$ , true so long as  $q > 2$ ). To show the claim, then, we fix an arbitrary  $j \in \{0, \dots, n\}$ , and consider an arbitrary set of coefficients  $(w_k)_{k < j}$  for which each  $w_k \in \{-1, 0, 1\}$  and at least one  $w_k$  is nonzero. It suffices to demonstrate, in this setting, that the sum  $\sum_{k < j} 2^{2+k} \cdot w_k$  cannot attain any of the values  $\{-2, -1, 0, -1, 2\}$ .

We first show that this property holds “over  $\mathbb{Z}$ ”; that is, we interpret all quantities (including the input bits  $v_{0,0}, \dots, v_{0,n-1}, v_{1,0}, \dots, v_{1,n-1}$ , as well as all scalars) as integers. We denote by  $k^*$  the smallest index  $k \in \{0, \dots, j-1\}$  for which  $w_{k^*} \neq 0$ . Clearly the sum  $\sum_{k=0}^{k^*} 2^{2+k} \cdot w_k \notin \{-2, -1, 0, -1, 2\}$ ; in fact, this sum equals either positive or negative  $2^{2+k^*}$  (whose absolute value is at least 4). In particular, this sum’s residue class modulo  $2^{3+k^*}$  is  $2^{2+k^*}$ .

We now argue that  $\sum_{k=0}^{j-1} 2^{2+k} \cdot w_k \notin \{-2, -1, 0, -1, 2\}$  as an integer. Informally, adding higher powers of two cannot change the sum’s residue class modulo  $2^{3+k^*}$ , which we have already seen equals  $2^{2+k^*}$ . More formally,

$$\sum_{k=0}^{j-1} 2^{2+k} \cdot w_k \equiv \sum_{k=0}^{k^*} 2^{2+k} \cdot w_k \equiv 2^{2+k^*} \pmod{2^{3+k^*}};$$

this follows from the equality  $\sum_{k=k^*+1}^{j-1} 2^{2+k} \cdot w_k \equiv 0 \pmod{2^{3+k^*}}$  (which actually holds regardless of the  $w_k$ ). As the above equation’s right-hand residue class modulo  $2^{3+k^*}$  is necessarily unequal to those of  $-2, -1, 0, 1$ , and  $2$ , the sum on its left-hand side is necessarily unequal to each of these values as an integer.

Finally, in light of the above, ensuring that  $\sum_{k=0}^{j-1} 2^{2+k} \cdot w_k \notin \{-2, -1, 0, -1, 2\}$  in  $\mathbb{F}_q$  amounts simply to showing that  $\sum_{k=0}^{j-1} 2^{2+k} \cdot w_k \in \{-q+3, \dots, q-3\}$  as an integer. In the worst case, this sum has an absolute value of  $\sum_{k=0}^{n-1} 2^{2+k} = 4 \cdot (2^n - 1)$ ; the statement of the first part of the theorem easily follows.

For the second part of the theorem, if  $v_i \leq v_{1-i}$  then there exists exactly one  $j \in \{0, \dots, n\}$  such that  $c_{i,j} = 0$ , and we have seen that this is the first  $j$  for which for every  $k < n$  it holds that  $v_{0,k} \neq v_{1,k}$  (in case  $v_i = v_{1-i}$ , we have that  $c_{i,n} = 0$ ). The algorithm shuffles the vector  $c_{i,0}, \dots, c_{i,n}$  and then multiplies each element with an independent random scalar. Thus, if  $v_i \leq v_{1-i}$  exactly one element would be 0, and all the rest are random in  $\mathbb{F}_q \setminus \{0\}$ . If  $v_i > v_{1-i}$ , then all elements in  $c_{i,0}, \dots, c_{i,n}$ , so after shuffling and multiplying with independent random scalars, the vector  $d_{i,0}, \dots, d_{i,n}$  distributes uniformly in  $\mathbb{F}_q^{n+1} \setminus \{0\}^{n+1}$ .  $\square$

## C Proof of Theorem 4.5

*Proof.* Regarding complexity, each party holds an input of length  $n$  bits. It secret shares each one of the bits using a field of size  $2^n$ , and therefore it sends  $O(n^2)$  bits. In addition, the parties run a secure coin-tossing protocol to toss  $\lambda$  coins, and its complexity is  $O(\lambda)$ . In total we get  $O(n^2 + \lambda)$ .

We consider the protocol in an idealized (i.e., a hybrid) model where  $\Pi_{\text{CT}}$  is a secure coin-tossing protocol. Replacing it with a secure protocol follows by a composition theorem. Since Functionality 4.1 is deterministic, we can separately show correctness and privacy, and we do not have to consider the joint distribution of the view of the adversary and the output of all parties. We start with showing correctness.

**Correctness.** We first analyze the security of a modified protocol in which the two parties run  $\pi_{\text{CT}}$  and receive all the randomness they need for running Algorithm 1. Note that this means that the parties do not use the pseudorandom generator at all throughout the execution.

All operations in Algorithm 1 on the values  $(v_{0,0}, \dots, v_{0,n-1}), (v_{1,0}, \dots, v_{1,n-1})$  are linear, and there-

fore working on shares of those vectors, applying Algorithm 1 on each one of the shares separately, and then reconstructing, results with the same output as running Algorithm 1 directly on and the values  $(v_{0,0}, \dots, v_{0,n-1}), (v_{1,0}, \dots, v_{1,n-1})$ .  $P^*$  then runs Algorithm 2 on the reconstructed values, and from Theorem 4.4, the result  $(b_0, b_1)$  is  $(v_0 \leq v_1, v_1 \leq v_0)$ , exactly as computed by the functionality.

When replacing the random string with the result of  $G(s)$ , we get that the output is computationally-indistinguishable from the output of the functionality.

**Privacy.** Here we separate to two different case. The first case is where  $P_i$  is corrupted for some  $i \in \{0, 1\}$ , and the second is when  $P^*$  is corrupted.

**Case I:  $P_i$  is corrupted.** The simulator receives the input  $(v_i)$  and the output  $(b_0, b_1)$  of the corrupted party,  $P_i$ . It has to generate the view of  $P_i$ . The view of  $P_i$  consists of the following:

1. The internal random coins of  $P_i$ .
2. The result of the coin-tossing protocol,  $s$ .
3. The random shares  $(\langle v_{1-i,j} \rangle_i^q)_{j=0}^{n-1}$ .
4. The output  $b_i$  as sent by  $P^*$ .

Clearly, all except for the last two items are independent random coins, which are easily simulatable. For the last item, the simulator receives them as input. Simulating the view of  $P_i$  follows easily.

**Case II:  $P^*$  is corrupted.** The simulator receives the input and output of  $P^*$ , which consists of the two bits  $(b_0, b_1) = (v_0 \leq v_1, v_1 \leq v_0)$ . The view of  $P^*$  consists of the shares  $(\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n$  for both  $i = 0$  and  $i = 1$ .  $P^*$  in the protocol then reconstructs for every  $j \in \{0, \dots, n\}$  the values  $d_{0,j}, d_{1,j}$  where

$$\begin{aligned} & (d_{0,0}, \dots, d_{0,n-1}), (d_{0,0}, \dots, d_{0,n-1}) \\ & = \text{ComparisonInitial}((v_{0,0}, \dots, v_{0,n-1}), \\ & \quad (v_{1,0}, \dots, v_{1,n-1})) . \end{aligned}$$

From Theorem 4.4 it follows that all the values  $d_{0,j}, d_{1,j}$  are uniform distributed in  $\mathbb{F}_q \setminus \{0\}$  unless for one position, which reveals whether  $v_0 \leq v_1$  and/or  $v_1 \leq v_0$ . In any case, simulating the values of  $d_{0,0}, \dots, d_{0,n}, d_{1,0}, \dots, d_{1,n}$  is easy. The simulator then generates additive shares for those values, as obtained from  $P_0$  and  $P_1$  (note that the shared randomness between  $P_0$  and  $P_1$  is not in the view of the adversary).

We remark that the view of  $P^*$  also consists of the entire communication between  $P_0$  and  $P_1$ , however, since it has no information on the private shared key between  $P_0$  and  $P_1$ , all encrypted information between the two parties is indistinguishable from random (as follows from the secrecy of the encryption scheme), and therefore simulating this communication is straightforward.  $\square$

## D Malicious Security of Protocol 4.3

**Computing minimum on committed inputs.** We slightly modify the protocol and the functionality, and this is the way we will use it to implement Functionalities 3.1 and 3.3. In particular, we assume that the server already holds commitments of the input values  $v_0, v_1$ , and the parties have the openings. Looking ahead, the clients first commit to their orders to the server; Later, when the server pairs two clients and we run the matching protocol, the bank provides the commitments to the counterparties. We have:

**FUNCTIONALITY D.1** ( $\mathcal{F}_{\min}^{\text{com}}$  – Computing minimum on committed inputs).

Consider two players,  $P_0$  and  $P_1$ , and a server  $P^*$ .

**The functionality:**

1. **In case of a corrupted  $P_i$ :** The honest  $P^*$  inputs commitments  $(V_0, V_1)$ , and the honest  $P_{1-i}$  sends  $(v_{1-i}, r_{1-i})$  such that  $V_{1-i} = \text{Com}(v_{1-i}; r_{1-i})$ . Send the corrupted  $P_i$  the commitments  $(V_0, V_1)$  and receive back  $(v_i, r_i)$ . Verify that  $V_i = \text{Com}(v_i; r_i)$ , and otherwise abort.
2. **In case of a corrupted  $P^*$ :** Receive from the each honest  $P_j$  the input  $(v_j, r_j)$ , and compute  $V_j = \text{Com}(v_j; r_j)$  for  $j \in \{0, 1\}$ . Send  $(V_0, V_1)$  to  $P^*$ .
3. Compute  $b_0 := (v_0 \leq v_1)$  and  $b_1 := (v_1 \leq v_0)$ .
4. Send  $b_0$  to  $P_0$ ,  $b_1$  to  $P_1$  and  $(b_0, b_1)$  to  $P^*$ . In addition, send to all parties  $\min\{v_0, v_1\}$ .

To implement this functionality, we slightly modify Protocol 4.3: The parties do not send the commitments in the first step. Instead,  $P^*$  sends both parties both commitments; In the last round, if a party outputs true, then it sends a new commitment  $V'_i$ , and proves using ComEq that the two commitments are the same, and opens  $V'_i$  to  $P^*$ .  $P^*$  verifies that the commitments are the same, and sends this information also to  $P_{1-i}$  as long as  $b_0 \neq b_1$  (i.e., if they are both 1, then both parties already know the minimum). Denote the modified protocol as  $\Pi_{\min}^{\text{com}}$ .

**Theorem D.2.**  $\Pi_{\min}^{\text{com}}$  securely computes the  $\mathcal{F}_{\min}^{\text{com}}$  (Functionality D.1) in the presence of a malicious adversary who corrupts a single party. Each party has to send or receive  $O(n(n + \lambda))$  bits where  $n$  is the length of the input and  $\lambda$  is the (computational) security parameter.

*Proof.* We prove the security of the protocol in the hybrid model, where the underlying functionalities are  $\mathcal{F}_{\text{ZK}}(\mathcal{R})$  (see Section A.1). Using the composition theorem, we later conclude security in the plain model.

We wrote the protocol in the plain model for brevity; We convert its description to a protocol in the  $(\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}}), \mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{BitProof}}), \mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{OneMany}}), \mathcal{F}_{\text{CT}})$ -hybrid model and prove that the protocol is secure in this model.

We discussed how to implement those functionalities in Section A.1), and using composition theorem we conclude that the protocol in the plain model is secure.

We need to prove that for any malicious adversary  $\mathcal{A}$ , the view generated by the simulator  $\mathcal{S}$  above is indistinguishable from the output in the real (hybrid)-model, namely:

$$\{\text{IDEAL}_{\mathcal{F}, \mathcal{S}(z), i}(u_0, u_1, \lambda)\} \approx_{\text{comp}} \{\text{REAL}_{\pi, \mathcal{A}(z), i}(u_0, x_1, \lambda)\}$$

We start with a simulator for the case where  $P_i$  is corrupted, for  $i \in \{0, 1\}$ . We later show a simulator for the case where  $P^*$  is corrupted. Without loss of generality, assume that  $i = 0$ .

### The simulator $\mathcal{S}$ :

1. The simulator invokes the adversary  $\mathcal{A}$  on an auxiliary input  $z$
2. The simulator receives from the trusted party the commitment  $V_0, V_1$ .
3. Simulate  $\mathcal{F}_{\text{CT}}$ : Choose a random  $s$  and give it to the adversary. If the adversary does not reply with OK, then send  $\perp$  to the trusted party.
4. Choose random shares for  $\langle v_{1,j} \rangle_0^q$  uniformly at random, and then compute commitments  $V_{1,j,0}$ .
5. Choose random commitments  $V_{1,j,1}$  such that  $\prod_{j=0}^{n-1} (V_{1,j,0} \cdot V_{1,j,1})^{2^{n-1-j}} = V_1$ .
6. Send  $\mathcal{A}$  the shares  $\langle v_{1,j} \rangle_0^q$ , the full array  $(V_{1,j,0}, V_{1,j,1})_{j \in N}$ , and open the commitments  $V_{1,j,0}$  for  $j \in N$ . Moreover, simulate  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}})$ ,  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{BitProof}})$  by sending accept to  $\mathcal{A}$ .
7. Receive from  $\mathcal{A}$  the full array  $(V_{0,j,0}, V_{0,j,1})_{j \in N}$  and the openings  $\langle v_{1,j} \rangle_1^q$ . Check that the openings indeed open  $V_{1,j,1}$ , for  $j \in N$ .
8. Simulate the  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}})$  invocation of Step 6: receive from  $\mathcal{A}$  as part of its invocation the witness for  $V_i$ , which consists of  $(v_i, r_i)$ , and the witness  $(v_i, r'_i)$  of  $\prod_{j=0}^{n-1} (V_{i,j,0} \cdot V_{i,j,1})^{2^{n-1-j}}$ . Verify that indeed  $V_i = \text{Com}(v_i; r_i)$  and that the  $\text{Com}(v_i, r'_i) = \prod_{j=0}^{n-1} (V_{i,j,0} \cdot V_{i,j,1})^{2^{n-1-j}}$ . If these conditions hold, return accept to Adv; otherwise, return reject.
9. Simulate the  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{BitProof}})$  invocation of Step 6 of each instance  $V_{i,j,0} \cdot V_{i,j,1}$  for every  $j \in N$ ; i.e., from each invocation also receive a witness and check that the witness satisfies the relation  $\mathcal{R}_{\text{BitProof}}$ . For each check return accept or reject.
10. If all checks pass, send  $v_i$  to the trusted party, where  $v_i$  was received when simulating the  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}})$  functionality in the previous steps. Otherwise, send  $\perp$ .
11. Receive  $(b_0, b_1, v)$  from the trusted party as  $P_i$ 's output.
12. Run Algorithm 1 on the shares provided to  $\mathcal{A}$  in the previous steps to obtain with randomness  $s$  (which was the result of the  $\mathcal{F}_{\text{CT}}$ ) to obtain  $(\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n$  and the randomnesses  $(s_{0,j,i})_{j=0}^n, (s_{1,j,i})_{j=0}^n$  and  $(\langle D_{0,j,1-i} \rangle_{j=0}^n, \langle D_{1,j,1-i} \rangle_{j=0}^n)$ .

13. Receive from  $\mathcal{A}$  the shares  $(\langle d_{0,j} \rangle_i^q)_{j=0}^n, (\langle d_{1,j} \rangle_i^q)_{j=0}^n$  and the randomnesses  $(s_{0,j,i})_{j=0}^n, (s_{1,j,i})_{j=0}^n$  as in Step 9. Verify that those are the same values as computed.
14. Compute  $D_{i,j} := \text{Com}(\langle d_{i,j} \rangle_i^q; s_{i,j,i}) \cdot D_{i,j,1-i}$ , for each  $j \in N$ , and then simulate invoking  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{OneMany}})$  while sending  $\mathcal{A}$  to output 1.
15. If  $b_1 = 1$ , then choose a random  $V'_1 = \text{Com}(v, r')$ . Send to  $P_0$  the commitment  $V'_1$ , as well as accept for the  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}})$  on the instance  $(V_1, V'_1)$ .
16. If  $b_0 = 1$ , then receive from  $\mathcal{A}$  commitment  $V'_0$ , and simulate an invocation of  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}})$  where expect to receive the opening of  $V_0$  and  $V'_0$ . If indeed the openings are correct and  $V'_0$  is a commitment of  $v_0$ , then send OK to the trusted party. Otherwise, send  $\perp$ .
17. The simulator outputs the view of the adversary in the above execution.

To show indistinguishability between the real and ideal execution, we consider the following experiments:

- Exp<sub>1</sub>: This is the real execution. The adversary  $\mathcal{A}$  is run with the honest  $P_1$  and  $P^*$ . The output of the experiment is the view of the adversary and the output of the honest parties  $P_1$  and  $P^*$ .
- Exp<sub>2</sub>: We run the simulator  $\mathcal{S}$  with the trusted party of Functionality D.1 with the following difference: When the honest party  $P_1$  sends  $v_1$  to the trusted party, the trusted party delivers  $v_1$  to  $\mathcal{S}$ . Then,  $\mathcal{S}$  uses  $v_1$  as the input of  $P_1$  to simulate the first steps of the protocol instead of choosing random commitments that match  $V_1$  as in the simulation. The output of this experiment is the output of all honest parties as determined by the trusted party and the output of the simulator.
- Exp<sub>3</sub>: this is just as the ideal execution: We run the simulator  $\mathcal{S}$  as prescribed above. The output of this experiment is the of all honest parties as determined by the trusted party and the output of the simulator.

We claim that the output of Exp<sub>1</sub> is computationally indistinguishable from Exp<sub>2</sub>. In particular, since the simulator receives the input of  $P_1$ , it essentially runs the honest  $P_1$  while exactly simulating all  $\mathcal{F}_{\text{ZK}}$  functionalities for the different relations in the protocol. The output of the protocol in the real execution is determined from the shares that  $P^*$  receives, which it then invokes Algorithm 2. In the simulation, the simulator uses the input  $v_0$  as extracted when invoking the  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{CompEq}})$  functionality, sending  $v_0$  to the trusted party, which then computes  $(b_0, b_1) = (v_0 \leq v_1, v_1 \leq v_0)$ . From the binding property of the commitment scheme, it is infeasible for the adversary to come up with a different  $v_0$  that does not match  $V_0$ . Correctness of the protocol from similar reasoning as in the semi-honest case, and from linearity of Algorithm 1.

We then claim that the output of Exp<sub>2</sub> is computationally indistinguishable from Exp<sub>3</sub>. The only difference is the view of the adversary are the commitments of  $V_{1,j,1}$  that are never

opened. Besides of those values, the entire view of the adversary is exactly the same (those are just accept messages received from the  $\mathcal{F}_{\text{ZK}}$  functionality). However, from the hiding property of the commitment scheme, the view of the adversary is indistinguishable.

**A corrupted  $P^*$ .** In case where  $P^*$  is corrupted, then the simulator simulates the entire communication between  $P_0$  and  $P_1$  as random strings.  $P^*$  has no input, and it receives from the trusted party the two bits  $(b_0, b_1)$ , the minimum  $M$ , and the commitments  $V_0$  and  $V_1$ . From the bits  $(b_0, b_1)$  it generates the two vectors  $(d_{0,0}, \dots, d_{0,n})$  and  $(d_{1,0}, \dots, d_{1,n})$  just as the simulator in the semi-honest case. It chooses random  $s_{0,j}, s_{1,j}$  and chooses random shares:

$$\begin{aligned} d_{0,j} &:= \langle d_{0,j} \rangle_0^q + \langle d_{0,j} \rangle_1^q, & s_{0,j} &:= s_{0,j,0} + s_{0,j,1} \\ d_{1,j} &:= \langle d_{1,j} \rangle_0^q + \langle d_{1,j} \rangle_1^q, & s_{1,j} &:= s_{1,j,0} + s_{1,j,1} \end{aligned}$$

It simulates  $P_0$  sending  $(\langle d_{0,j} \rangle_0^q)_{j=0}^n, (\langle d_{1,j} \rangle_0^q)_{j=0}^n$  and the randomnesses  $(s_{0,j,0})_{j=0}^n, (s_{1,j,0})_{j=0}^n$  to  $\mathcal{A}$ , and  $P_1$  sending  $(\langle d_{0,j} \rangle_1^q)_{j=0}^n, (\langle d_{1,j} \rangle_1^q)_{j=0}^n$  and the randomnesses  $(s_{0,j,1})_{j=0}^n, (s_{1,j,1})_{j=0}^n$  to  $\mathcal{A}$ . It then simulates the invocation of  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{OneMany}})$ : If  $\mathcal{A}$  does not provide the correct statement,  $(D_{i,j})_{j=0}^n$  for  $D_{i,j} := \text{Com}(d_{i,j}; s_{i,j})$ , and the correct witness (consisting of the index of the non-zero coordinate in  $d$  and the opening of its commitment), then return reject to  $\mathcal{A}$  and send  $\perp$  to the trusted party. Otherwise, send accept to  $\mathcal{A}$  and send OK to the trusted party. At the end of the invocation, if some party  $P_j$  has to open the commitment to  $\mathcal{A}$ , then it sends a commitment  $V_j' = \text{Com}(M; r)$  with some random  $r$ , and simulates  $\mathcal{F}_{\text{ZK}}(\mathcal{R}_{\text{ComEq}})$  on instance  $V_j, V_j'$  as accept.

It is easy to see that the simulator perfectly simulates the view of the adversary in the real execution. This follows from the properties of our comparison algorithm (Theorem 4.4), and that the simulator perfectly simulates the  $\mathcal{F}_{\text{ZK}}$  functionality. Finally, any wrong information  $\mathcal{A}$  sends to  $\mathcal{F}_{\text{ZK}}$  results in false in the real execution. In contrast, in the ideal execution, the simulator sends  $\perp$  to the trusted party which results in  $\perp$  to both  $P_0, P_1$  in the ideal execution.  $\square$

**Computing  $\mathcal{F}_{\text{C2C}}$  functionality (Functionality 3.2.)** To conclude client-to-client matching, the parties invoke  $2|U|$  times  $\mathcal{F}_{\text{min}}$ . For each possible symbol in  $U$ , the parties first commit to their short and long exposure for that symbol. They then invoke the committed minimum functionality twice: Once when  $P_0$  inputs its long exposure and  $P_1$  inputs short, and vice-versa. If a party is not interested in some symbol or in a particular side then it simply inputs 0. The bank receives as output all the matches and can execute them directly. Security follows from just the composition theorem. We have:

**Corollary D.3.** *The above protocol securely computes Functionality 3.2 in the presence of one malicious client or a malicious server.*

## E Bank-to-Client Matching

In this section, we present a two-party variant of our main three-party comparison protocol which can be used for bank-to-client matching. We describe the functionality and then the protocol.

**FUNCTIONALITY E.1** (Two-party min,  $\mathcal{F}_{\text{min}}^{\text{B2C}}$ ).

**Input:** The server  $P_0$  holds  $v_0$ , and  $P_1$  holds  $v_1$ , both in  $\{0, \dots, 2^n - 1\}$ .

**Output:** Both parties receive  $\min\{v_0, v_1\}$ .

**PROTOCOL E.2** (Two-party protocol  $\Pi_{\text{min}}^{\text{B2C}}$ ).

**Input:**  $P_0$  and  $P_1$  hold integers  $v_0$  and  $v_1$ , respectively, in  $\{0, \dots, 2^n - 1\}$ .

**Setup phase:** A commitment scheme (Com) and an encryption scheme (Gen, Enc, Dec) are chosen.  $P_0$  runs  $(pk, sk) \leftarrow \text{Gen}(1^\lambda)$ , and authenticates its public key  $pk$  with  $P_1$ .

**The protocol:**

1.  $P_0$  proceeds as follows: (**Encryption**):

- (a) Commit  $V_0 \leftarrow \text{Com}(v_0; r_0)$ , and send  $V_0$  to  $P_1$ .
- (b) Compute the bit decomposition  $v_0 = \sum_{j \in N} 2^{n-1-j} \cdot v_{0,j}$ , for bits  $v_{0,j} \in \{0, 1\}$ .
- (c) For each  $j \in N$ : compute additive homomorphic encryptions  $A_j = \text{Enc}_{pk}(v_{0,j})$ .
- (d) Send the full array  $(A_j)_{j=0}^{N-1}$  to  $P_1$ .
- (e) Compute:

- i.  $\pi \leftarrow \text{ComEq.Prove}(V_0, \prod_{j=0}^{N-1} (A_j)^{2^j})$ ,

- ii.  $\pi_j \leftarrow \text{BitProof.Prove}(A_j)$ , for all  $j \in N$ ,

$P_0$  sends  $\pi$ , and  $(\pi_j)_{j=0}^{N-1}$  to  $P_1$ .

2.  $P_1$  proceeds as follows: (**Computing the minimum**):

Receive  $(\pi, (\pi_j)_{j=0}^{N-1})$  and verify the following:

- (a)  $\text{ComEq.Verify}(\pi, V_0, \prod_{j=0}^{N-1} (A_j)^{2^j})$ ,
- (b)  $\text{BitProof.Verify}(\pi_j, A_j)$  for each  $j \in N$ . If any of these checks fail,  $P_1$  aborts.
- (c) Run Algorithm 1, in parallel, on the ciphertexts  $A_j$  and on its own secret inputs. That is,  $P_1$  runs:

$$\begin{aligned} & \left( (D_{0,j})_{j=0}^N, (D_{1,j})_{j=0}^N \right) \leftarrow \\ & \text{ComparisonInitial} \left( (A_j)_{j=0}^{N-1}, (v_{1,j})_{j=0}^{N-1} \right) \end{aligned}$$

recall that the algorithm shuffles the two result vectors inside.

- (d)  $P_1$  sends  $((D_{0,j})_{j=0}^N, (D_{1,j})_{j=0}^N)$  to  $P_0$ .

### 3. Party $P_0$ (Output reconstruction):

- (a) Decrypt  $d_{i,j} = \text{dec}_{sk}(D_{i,j})$  for each  $i = \{0, 1\}$  and  $j = \{0, \dots, n-1\}$
- (b) Execute Algorithm 2, that is:

$$(b_0, b_1) := \text{ComparisonFinal} \left( (d_{0,j})_{j=0}^N, (d_{1,j})_{j=0}^N \right)$$

- (c) Write  $u$  for the index such that  $b_u$  is true and sets  $u := 0$  if both are.
- (d) Compute  $\pi' \leftarrow \text{OneMany.Prove} \left( (D_{u,j})_{j=0}^{n-1} \right)$ , and sends  $\pi'$  and  $u$  to  $P_1$ . If  $u = 1$  then send also  $v_1$ .

### 4. Party $P_1$ (output reconstruction):

- (a)  $P_1$  verifies  $\text{OneMany.Verify} \left( \pi', (D_{u,j})_{j=0}^{n-1} \right)$ . If verification passes, then  $P_1$  outputs  $v_u$ .

**Theorem E.3.** *Protocol E.2 securely computes Functionality E.1 in the presence of a malicious  $P_0$  or a semi-honest  $P_1$ , assuming secure commitments and zero-knowledge functionalities.*

*Proof.* Since the protocol is not symmetric, we separate between corrupted  $P_0$  and corrupted  $P_1$ .

**Security against a malicious  $P_0$ :** We show security assuming an ideal commitment and a zero-knowledge scheme. The simulator plays the role of the ideal functionalities for those sub-protocols and receives in particular the input  $v_0$  and the decomposition bits when  $P_0$  calls to the zero-knowledge and commitment functionalities. It verifies that  $P_0$  provided consistent input. If not, the simulator aborts in the corresponding functionality.

If the inputs are consistent, then given the input  $v_0$  it sends it to the ideal functionality and receives back the minimum and whether the two integers are the same. It computes from that  $(b_0, b_1)$ , and simulates values  $d_{0,j}, d_{1,j}$  uniformly at random as per Theorem 4.4, and sends to  $P_1$  encryption of those values under  $P_0$ 's public key. It then verifies that the one-out-of-many proof of  $P_0$  is correct, and if so it sends OK to the trusted party, in which it delivers to  $P_1$  its output in the ideal world. It is easy to see that the simulator perfectly simulates the view of the adversary in the ideal world.

**Security against a semi-honest  $P_1$ :** In that case, the simulator sends to the adversary a commitment to 0, and bit decomposition as all zeros. Moreover, it simulates the zero-knowledge proofs as succeeding. It then receives encryptions of  $(d_{0,j}, d_{1,j})$ . It receives from the trusted party the minimum value and simulates the final ZK proof as successful.  $\square$

**Concrete instantiation.** We implement Protocol E.2 using the ElGamal encryption scheme, where the message  $m$  is part of the exponent (thereby we get additive homomorphism). I.e., for a public key  $h \in \mathbb{G}$ ,  $\text{Enc}_{pk}(m) = (g^r, h^r g^m)$ . Note that

$P_1$  does not have to decrypt the ciphertexts, but just identify an encryption of 0. Given a secret key  $x$  such that  $h = g^x$  and a ciphertext  $c = (c_1, c_2)$ , this is done by simply comparing  $c_2/c_1^x$  to  $g^0$ .

**Implementing Functionality 3.1.** To implement the bank-to-client functionality, the parties invoke  $2|U|$  times the two-party minimum functionality. For each possible symbol in  $U$ , the parties invoke the minimum functionality where once the bank inputs its interest in symbols to buy and the client in sell, and one time where the bank inputs whether it wishes to sell and the client inputs whether it wishes to buy. If a party is not interested in some symbol or in a particular side, it simply inputs 0. The results of all minimum invocations will reveal the matches. Security follows from composition. We have:

**Corollary E.4.** *The above protocol securely computes Functionality 3.1 ( $\mathcal{F}_{B2C}$ ) in the presence of a semi-honest bank or a malicious client.*

## F Realizing Functionality 3.3

In this Section we present protocol F.1 for the functionality 3.3 in the  $\mathcal{F}_{\min}^{\text{com}}$ -hybrid (Functionality D.1) model. When matching begins, orders pertaining to the same security and of opposite direction are matched. After each particular match, the server decrements both orders' quantities by the matched amount, and dequeues whichever among the two orders is empty (necessarily at least one will be). The commitments' homomorphic property allows the engine to appropriately decrement registrations between matches.

### PROTOCOL F.1 ( $\Pi_{MC}$ —multiparty matching).

Upon initialization,  $P^*$  initializes a list  $P = \emptyset$  and two vectors  $\mathcal{L}$  and  $\mathcal{S}$  of size  $n$ , where  $n$  bounds the number of possible clients.

$\Pi_{MC}.\text{Register}(P_i)$ : When the command is invoked, then  $P_i$  sends to  $P^*$  commitments  $\tilde{L}_i = \text{Com}(L_i; r_i)$  and  $\tilde{S}_i = \text{Com}(S_i; s_i)$  for some random  $r_i, s_i$ .  $P^*$  set  $\mathcal{L}[i] = \tilde{L}_i$  and  $\mathcal{S}[i] = \tilde{S}_i$ . The party  $P_i$  locally stores  $(L_i, r_i)$  and  $(S_i, s_i)$ .  $P^*$  also adds  $i$  to  $P$ .

$\Pi_{MC}.\text{Process}()$ :

1.  $P^*$  chooses a random ordering  $O$  over all pairs of  $P$ .
2. For the next pair  $(i, j) \in O$ , try to match between  $P_i$  and  $P_j$ :
  - (a) Invoke Functionality D.1 where  $P_i$  inputs  $(L_i, r_i)$  and  $P_j$  inputs  $(S_j, s_j)$ ,  $P^*$  inputs  $(\mathcal{L}[i], \mathcal{S}[j])$ . Let  $M_0$  be the resulting minimum.  $P_i$  sets  $L_i = L_i - M_0$  and  $P_j$  sets  $S_j = S_j - M_0$ . If  $M_0 \neq 0$ , then execute the match.  $P^*$  homomorphically evaluates the commitments  $(\mathcal{L}[i], \mathcal{S}[j])$  to subtract  $M_0$ .

- (b) Invoke Functionality [D.1](#) where  $P_i$  inputs  $(S_i, s_i)$  and  $P_j$  inputs  $(L_j, r_j)$ ,  $P^*$  inputs  $(S[i], \mathcal{L}[j])$ . Let  $M_1$  be the resulting minimum.  $P_i$  sets  $S_i = S_i - M_1$  and  $P_j$  sets  $L_j = L_j - M_1$ . If  $M_1 \neq 0$ , then execute the match.  $P^*$  homomorphically evaluates the commitments  $(\mathcal{L}[i], S[j])$  to subtract  $M_1$ .

**Theorem F.2.** *Protocol [F.1](#) securely implements Functionality [3.3](#) in the  $\mathcal{F}_{\min}^{\text{com}}$ -hybrid model (Functionality [D.1](#)) in the presence of a malicious client or a semi-honest server.*

*Proof.* We first simulate the case where some client is malicious:

1. Invoke the adversary  $\mathcal{A}$ . Whenever the adversary sends  $\tilde{L}_i, \tilde{S}_i$ , store  $\tilde{L}_i$  and  $\tilde{S}_i$  in  $\mathcal{L}[i]$  and  $S[i]$ , respectively.
2. Whenever received  $(i, j, M_0, M_1, b_i)$  from the trusted party, choose random commitments  $\mathcal{L}[j]$  and  $S[j]$  and simulate two invocations of  $\mathcal{F}_{\min}^{\text{com}}$ :
  - (a) Send to  $P_i$  the commitment  $\mathcal{L}[i]$  and a random commitment  $S[j]$ . Receive back  $b_i$  and  $M_0$ . Homomorphically subtract  $M_0$  from  $\mathcal{L}[i]$ .
  - (b) Send to  $P_i$  the commitment  $S[i]$  and a random commitment  $\mathcal{L}[j]$ . Receive back  $b_j$  and  $M_1$ . Homomorphically subtract  $M_1$  from  $\mathcal{L}[i]$ .

From inspection, it is easy to see that the view of  $P_i$  is the same in the simulation and in the real world since the simulator perfectly simulates the invocations of Functionality [D.1](#) from the information it receives from the trusted party. Moreover, the underlying committed values of  $P^*$  in the real execution are the same as in the ideal process for the same ordering  $O$  (which have the exact same distribution in the real and ideal). This can be shown by induction over the different invocations of the ordering of  $O$ : in each execution, the functionality [D.1](#) gives  $P^*$  the same as in the ideal execution of  $\mathcal{F}_{\text{MC}}$ . From the homomorphic evaluation of the commitment scheme, the subtraction of the values is equivalent to  $\mathcal{F}_{\text{MC}}$  subtracting the minimum values.

**The case of a corrupted  $P^*$ .** We show security in the presence of a *semi-honest*  $P^*$ . First, whenever received a message registered( $P_i$ ) from the trusted party, simulate  $P_i$  sending to  $P^*$  some random commitments  $\tilde{L}_i, \tilde{S}_i$ , and store those commitments.

The simulator then repeatedly receives from the trusted party values  $(i, j, M_0, M_1, b_0, b_1)$  and has to simulate the view of  $P^*$  in the two invocations of  $\mathcal{F}_{\min}^{\text{com}}$ . First, from all those values, it can extract random coins that will result in the same random ordering  $O$  as the one that the trusted party chose. Moreover, for each such  $(i, j, M_0, M_1)$ ,  $(b_0^0, b_0^1, b_1^0, b_1^1)$ , the simulator pretend sending the adversary  $\tilde{L}_i, \tilde{S}_j$  as coming from  $\mathcal{F}_{\min}^{\text{com}}$ ; and then the output  $(b_0^0, b_0^1)$  and  $M_0$ . For the second invocation, pretend sending the adversary  $\tilde{S}_i, \tilde{L}_j$  and the output

$(b_1^0, b_1^1)$  and  $M_1$ . Then, update  $\tilde{L}_i, \tilde{S}_i$  and  $\tilde{L}_j, \tilde{S}_j$  as in the protocol. From inspection, it is clear that the view of the adversary is the same in both executions.  $\square$

## G Range Bank-to-Client Functionality

In this section, we consider the setting in which the client submits a minimum and a maximum amount per stock instead of a single amount. The current system is set to accept a minimum and a maximum quantity from the clients to be matched against the bank's inventory. As a first step, the bank runs Functionality [3.1](#) on every client one by one only on their minimum quantity until the bank exhausts its inventory. For the case where the bank did not exhaust its inventory after processing all the clients, it runs again Functionality [3.1](#) on every client one by one on their maximum quantity. This process is presented in Functionality [G.1](#). The requirement on the minimum and maximum quantity is imposed by the business for this use case as a greedy approach in an effort to satisfy more clients with the bank's inventory. This is not a necessary requirement, but the bank views it as an additional feature to satisfy more clients. This greedy approach is not ideal since, for instance, the minimum quantity of the first client can exhaust the full inventory of the bank. We leave it as an open problem to devise a secure optimization algorithm for better allocations across all clients.

This additional range feature is considered only when we seek for matches between the bank and a client in an effort to maximize the number of matches the bank can accommodate from its own inventory. That said, we do not explicitly consider it for the client-to-client matching.

### FUNCTIONALITY G.1 ( $\mathcal{F}_{\text{B2C}}$ –Range Bank-to-client functionality).

The functionality is parameterized by the set of all possible securities to trade, a set  $U$ .

**Input:** The bank  $P^*$  inputs lists of orders  $(\text{symp}_i^*, \text{side}_i^*, \text{amount}_i^*)$  where  $\text{symp}_i \subseteq U$  is the security,  $\text{side}_i^* \in \{\text{buy}, \text{sell}\}$  and  $\text{amount}_i^*$  is an integer. A client  $P_i$  sends its list of the same format but with a minimum and a maximum amount,  $(\text{symp}_i^C, \text{side}_i^C, \text{MinAmount}_i^C, \text{MaxAmount}_i^C)$ .

**Output:** Initialize a list of Matches  $M$ . Choose a random order  $O$  over all clients.

- For the next client  $j$  in  $O$  and for every  $i$ , such that  $\text{symp}_i^* = \text{symp}_j^C$  and  $\text{side}_i^* \neq \text{side}_j^C$ , add  $(\text{symp}_i^*, \text{side}_i^*, \text{side}_j^C, v = \min\{\text{amount}_i^*, \text{MinAmount}_j^C\})$  to  $M$  and update  $\text{amount}_i^* = \text{amount}_i^* - v$ .

Then, check maximum amounts for the  $j$  in  $M$ :

- For the next  $j$  in  $O$  and every  $i$  such that  $\text{symb}_i^* = \text{symb}_j^C$  and  $\text{side}_i^* \neq \text{side}_j^C$ , add  $(\text{symb}_i^*, \text{side}_i^*, \text{side}_j^C, v = \min\{\text{amount}_i^*, \text{MaxAmount}_j^C\})$  to  $M$  and update  $\text{amount}_i^* = \text{amount}_i^* - v$ .

Implementing this protocol is quite straightforward, given the protocols we already have. The bank simply runs the minimum functionality against (random) ordering of the clients with their minimum amount, and then with their maximum amount, while updating the amount of each symbol along the way.

## H Prime Match System Performance Cont.

In Figure 7, we present a sample of the user interface of the Prime Match system.

## I Comparison to Generic MPC

In this section, we compare the performance of our protocol to generic MPC protocols. Concretely, we look at the client-to-client matching. Towards that end, we estimate the size of the circuit that we have to use to implement such a functionality.

Our main client-to-client matching protocol has two phases, in which the client first commits to their orders, and then the parties run in a second phase that the orders they provide match the committed ones. Implementing that using generic MPC protocols is relatively expensive as they must involve cryptographic primitives. Implementing group operations in a circuit, which involves exponentiation and public-key primitives is clearly prohibitively expensive, so we would like to use symmetric-key primitives.

Specifically, we look at a commitment scheme in which the server (the bank) chooses some random seed  $s$  and publishes it to all parties. To commit to a value  $v$ , a party chooses a random  $r$  and computes  $H(s, v, r)$  where  $H$  is a CRHF. This commitment scheme is secure when modeling  $H$  as a random oracle. When instantiating  $H$  using SHA256, this (heuristically) gives 128-bit security for the binding property (a collision implies breaking the commitment).

A Boolean circuit for SHA256 has 22K number of AND gates, and depth 1607 (see e.g., [1, 12]). This means that even ignoring the costs for comparisons themselves, we have:

- Using any garbling-like protocol (say, variants of BMR) and with 128-bit keys for the labels, a circuit that checks that the commitments are correct will require communicating  $22K \cdot 128 \cdot 2 \cdot 2$  bits (128 bit keys; using half gate optimization, so two ciphertexts per gate; for each symbol, we need at least two commitments - one from each one of the clients). The communication would be at least 722KB for processing a symbol. This is compared to roughly 25 KB per symbol by our protocol, as

reflected by Table 2. This is at least  $\times 30$  more expensive than our solution, while the actual performance is likely to be much more expensive. We ignored here the cost of the comparisons themselves, or if the clients have to be paired later again then the circuit has to compute the modified commitment after reducing the minimum (which requires further evaluations of SHA inside the circuit).

- Using any GMW-like protocol that requires number of rounds that is proportional to the depth of the circuit, the number of rounds would be at least 1607. This is compared to 3 rounds as in our protocols. This is an increase in a factor of  $\times 500$ .

## J Challenges in Implementation

Privacy-preserving auctions have been the holy grail of practical MPC since [9]. However, no financial institutions took a step forward to use such tools. To begin with, we got connected to the bank’s business department of quantitative research to test the appetite for a form of privacy-preserving auctions. Thus, the Axe inventory use case was chosen.

**Pre-Production Challenges.** Releasing a secure multiparty computation (MPC) product in a big organization that had no other products utilizing these techniques was itself challenging. Furthermore, there were no other products across the street (market) to showcase the feasibility of such a solution.

As a first step, a proof of concept (POC) was implemented to show its feasibility.

Given the innovative nature of the product, the green light for production was given after a long process. More specifically, the organization decided to test the appetite for the possible product of its valuable clients (hedge funds). Almost 20 hedge funds were visited in 6 months to demo the POC and to hear their highly appreciated feedback. Given the client feedback, the organization decided to move another step forward and ask the clients under what conditions they will utilize such a product. In particular, the clients provided a set of different features and requirements they would like to see in the product. Notable requirements were:

- No communication with other clients, only communication with the bank.
- No resources for preprocessing data.
- No installation of code on client’s machines - a web-based application was required.
- Stronger security guarantees than semi-honest security.
- Peer review of the solution.

The POC was updated to satisfy the above requirements. Another round of demos to clients was conducted. After several months the business gave the green light to build the product and allocated resources to enable it. The team consisted of cryptographers, quants, tech quants, expert programmers on

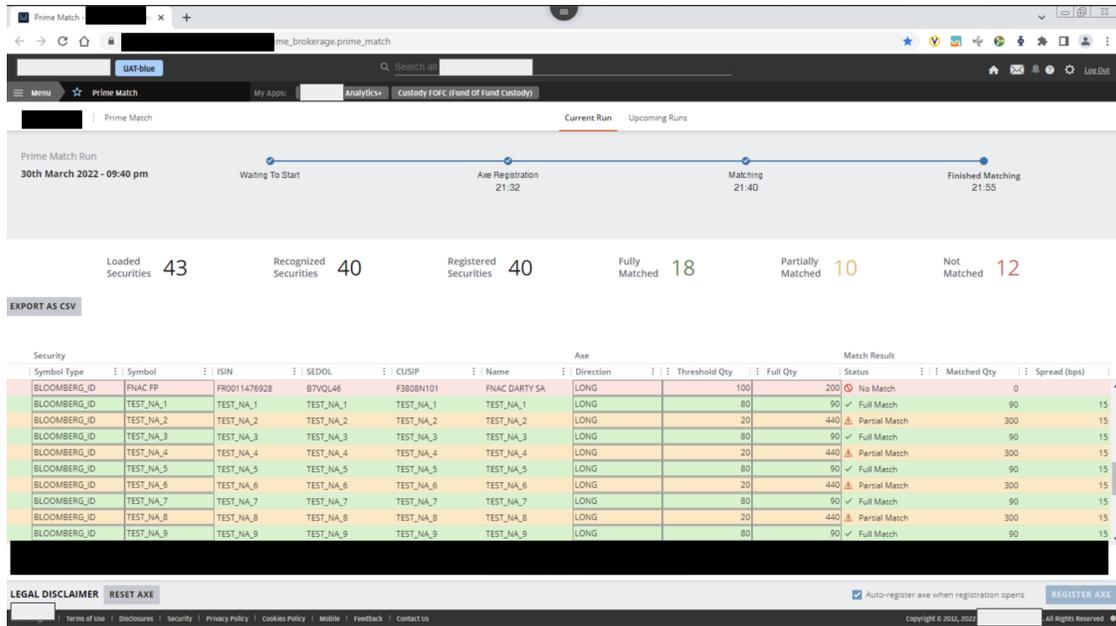


Figure 5: Prime Match User Interface of a completed matching run.

different topics (such as secure multiparty computation, UI, web-based development etc.), cybersecurity experts, product managers, legal, and stakeholders.

**Production Challenges - Lessons Learned** During the development of Prime Match, we faced several critical technical challenges. Those challenges were mainly due to the complex design of the organization’s infrastructure for security measures.

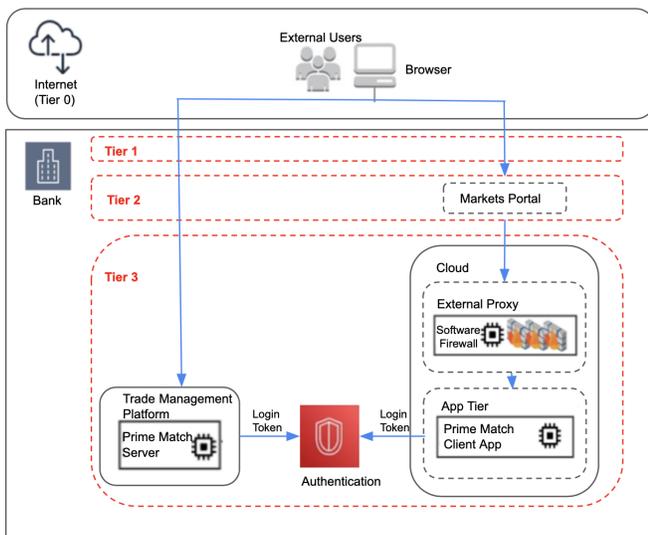


Figure 6: Prime Match system architecture connecting external users/clients to the bank’s network.

- Coding in the Trade Management platform and the Markets Portal, see Figure 6, needed to follow certain practices (which easy maintenance too) which were not followed in the POC stage. POC was reprogrammed to follow current practices.
- Only a limited set of open-source libraries are allowed in the Trade Management platform. For example, importing the open-source library emscripten docker image to the Trade Management platform. This process went through rigorous request and approval processes to show that this library does not contain any code that would secretly send data to the internet.
- Integration to the markets portal which is an existing service for external clients, in order to access some of the organization’s trading applications, was not easy since no other application required interaction between the server and the clients. Prior apps involve simple downloads of data from clients.
- Set up a pathway that allows network traffic to flow from the clients to the Prime Match server with the use of web socket protocol. In particular, setting up PaaS service in Tier 1 and 2, which stands for Perimeter Security as

Number of Symbols	Latency (sec)	Throughput (transactions/sec)	Registration received msg Size (MB)	Registration sent msg Size (MB)	Matching received msg Size (MB)	Matching sent msg Size (MB)
100	23.85 ( $\pm 9.19$ )	8.38	0.062	0.036	16.225	4.510
200	29.49 ( $\pm 2.85$ )	13.56	0.124	0.073	32.451	9.021
500	65.04 ( $\pm 1.41$ )	15.37	0.310	0.184	81.128	22.552
1000	136.93 ( $\pm 3.32$ )	14.60	0.621	0.370	162.257	45.105
2000	244.37 ( $\pm 10.31$ )	16.36	1.243	0.749	324.507	90.210
3000	428.48 ( $\pm 32.19$ )	14.00	1.865	1.128	486.739	135.260
4000	1077.60 ( $\pm 27.46$ )	7.42	2.487	1.507	647.758	180.043
5000	1314.44 ( $\pm 39.15$ )	7.60	3.109	1.886	809.191	224.641
6000	1558.58 ( $\pm 46.96$ )	7.69	3.731	2.265	971.875	269.997

Table 3: Performance of Bank-to-Client Prime Match for two clients. The number of symbols processed in the production code is double since clients provide both a minimum and a maximum quantity to be matched. This is reflected at the Throughput column which is multiplied by a factor of two.

a Service. Extensive review and approval process were required.

- Placing a physical server and setting up an Apache WS tunnel in tier 1 to enable external clients to transact and transact within tight latency requirements. In particular, we had to implement the entire WS tunnel proxy from scratch since the Web Sockets protocol was not used within these platforms before.
- Run several internal tests ranging from UI requirements to cyber security requirements (such as timing attacks).
- Extensive code review from internal and external experts.
- Integrate the UI component with the markets portal

which required extensive approvals.

- Integration of our code with WebAssembly to achieve near-native performance. To this end, we also had to handle garbage collection. See more details on the chosen programming languages in Section 5.

Overall each step required extensive review and approval processes. An important aspect that was not listed above is the last process of extensive testing, verification of the code, and beta testing the product with the help of a limited number of clients.

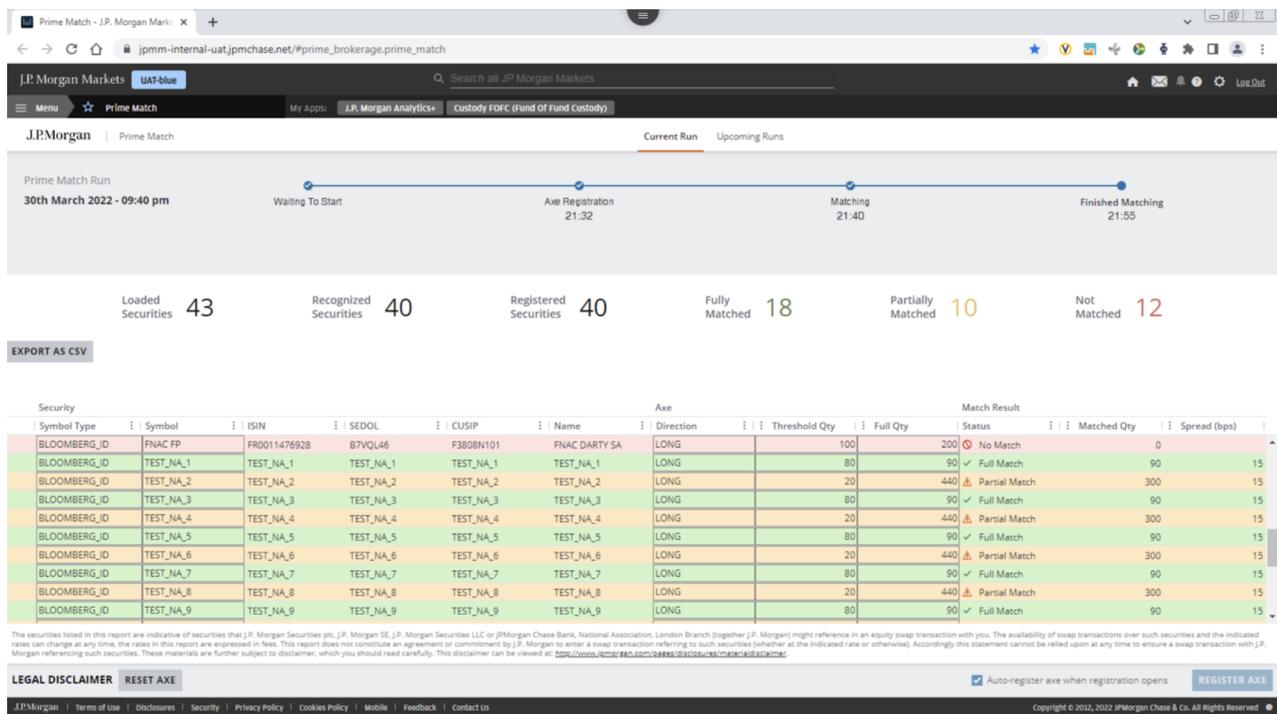


Figure 7: Prime Match User Interface of a completed matching run.