Privacy-Preserving Dark Pools

Extended Abstract

Gilad Asharov*
Bar Ilan University
gilad.asharov@biu.ac.il

Tucker Balch
J.P. Morgan AI Research
tucker.balch@jpmorgan.com

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

Manuela Veloso
J.P. Morgan AI Research
manuela.veloso@jpmorgan.com

ABSTRACT

A dark pool is a private forum for trading financial instruments such as equities and derivatives. The service is offered to traders, often representing large financial institutions, who aim to make large trades without telegraphing their intentions in a public venue such as the New York Stock Exchange or NASDAQ. Internally, the dark pool operates just like a public exchange, but the order book, a list of offers to buy and sell, is not visible to any of the participants. Only when matching orders are found the trade is executed and reported externally. Because the large orders are not visible to others, the risk of significant price moves is reduced.

The operators of dark pools are trustworthy, but traders are notoriously suspicious. Traders are concerned about any "leakage" of information that might alert others to their intentions and cause price moves against them. The operator sees all orders posted by all participants, including sensitive information such as the volume of the order and the bid or ask price. Because the operator of the dark pool is usually a large financial institution with its own investments in the assets being traded, it may have conflicts of interest with the clients it serves in the dark pool. While such conflicts are mitigated by law and regulation, they continue to exist.

We propose a new mechanism that reduce the need of traders to trust the operator of the dark pool. By adopting cryptographic techniques, we achieve a fully private and secure marketplace where sellers and buyers interact with the operator without exposing the volumes or prices of their orders to the operator or to any other counter party. These values remain secret until a matching order is found, and only afterward it will be revealed publicly.

KEYWORDS

Secure computation; Auctions

ACM Reference Format:
Gilad Asharov, Tucker Balch, Antigoni Polychroniadou, and Manuela Veloso. 2020. Privacy-Preserving Dark Pools. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 PROBLEM STATEMENT

Public exchanges such as the New York Stock exchange or NASDAQ act as auctioneers in a public double auction process. Potential buyers and sellers interact with the auctioneer, and submit their bids or asking prices. Buyers submit the maximal price they are willing to pay for the security, while sellers submit the minimal price they are willing to receive for selling the security. Both the seller and the buyer also submit the number of shares they are willing to trade. The auctioneer looks for matches between all orders it receives, and once matching orders are found, it executes the matching orders.

All orders are submitted to the public exchange in the clear, which inevitably impacts the market if the order is large: a seller that wishes to sell a large amount of securities and submits such an order to the exchange will immediately move the price downwards and against him, due to the sharp increase of the supply of that security. Similarly, a buyer that submits a bid for a large quantity of a particular stock will immediately move the price of the stock upwards and against him. Splitting the order into multiple smaller orders does not reduce market impact.

In order to decrease such market impacts, traders use dark pools. A dark pool is a private forum for trading financial products. Buyers and sellers send private orders that are visible only to the operator of the dark pool, and the operator executes an order matching system just like the public exchange. Only when matching orders are found, the orders are executed and reported in the public exchange. Because the large orders are not visible to others, the risk of significant price moves is reduced. There are more than 40 dark pools registered in the U.S., and around 14% of U.S. equity volume is being executed via dark pools [1]. The majority of the dark pools are being operated by investment banks. Yet, we will address in this paper one major weakness that exists in dark pools.

Challenge: Mistrust and conflicts of interest. Mistrust and conflict of interests exist between the operator of the dark pool and its participants. The operator sees all orders posted by all participants, including sensitive information such as the volume of the order and the bid or ask price. Usually, the operator itself is a large financial institution with its own investments in the assets being traded, and thus it may have conflicts of interest with the clients it servers in the dark pool. While such conflicts are mitigated by law and regulation, they continue to exist. Besides price moves, dark pool violations include also *front running*. Consider the situation where two clients submit two orders to the dark pool at the same time: Client A submits a sell order while the minimum price that he/she is willing to accept is $100 per share. Client B submits a buy order with the maximum price that he/she is willing to buy is $105. Clearly, we have a match. In that case, the price in which the trade will be executed is either the mid-price (102.5$), or according to the price of the order that was submitted first while the operator has

[*] The author contributed while employed at J.P. Morgan AI Research

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

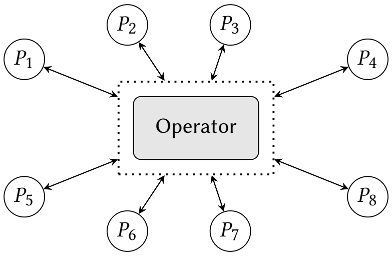

Figure 1: Privacy-preserving dark pool with a single operator. The orders sent to the operator are encrypted and the operator computes on the encrypted orders.

to specify such rules in advance. However, a malicious operator, or any other agent that was exposed to these two orders can then buy from Client A for $100 per share and sell to Client B for $105 per share. Despite law and regulation, the fact that such strategies and attacks exist increase the tension between the clients and the operator of the dark-pool.

2 OUR SOLUTION

Using advanced cryptographic techniques, we present a mechanism that addresses the main weakness of today's dark pool systems. Our solution is based on *Secure Multiparty Computation* (MPC) [2, 4, 6, 8, 9] and Fully-Homomorphic Encryption (FHE) techniques [5]. MPC allows a set of mutually distrustful parties to compute a joint function without revealing any information about their inputs besides what is being revealed by the output of the computation. FHE is a form of encryption that allows computation on encrypted data, generating an encrypted result which, when decrypted, equals to the result of the computation on the original data. We devise the following mechanism:

Privacy-preserving dark pool. A privacy preserving dark pool is a system where parties send their buy and sell orders (see Figure 1) to an operator in an encrypted way, which allows the operator to compute on the encrypted data whether there is a match between orders in the order book or not. The operator of the dark pool receives only encrypted orders – it cannot tell whether the order is Buy or Sell, the amount and the bid or ask price. When matching orders are found, information regarding the matching orders is being decrypted, and the operator can execute the matching orders in the public market. Since orders remain private to the operator, privacy-preserving dark pools have the potential to further reduce market impacts compared to standard dark pools, as well as to significantly reduce the tension between the participants of the dark pool and its operator.

Paper contributions. In the extended version of this paper, we present the following contributions:

- We identify common drawbacks in recent dark pools and address them using cryptographic techniques.
- We propose secure protocols for realizing the mechanism. We describe a protocol for realizing a privacy-preserving dark pool.
- We implement the mechanism and empirically evaluate it.

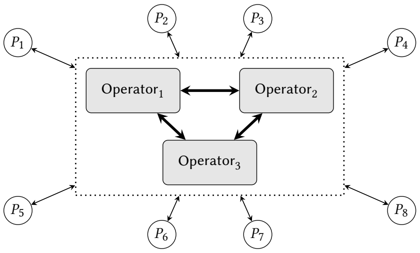

Figure 2: Distributed privacy-preserving dark pool. Three different operators, each runs its independent privacy-preserving dark pool.

Related work. A recent work by Catlidge, Smart and Talibi [3] considers dark pools in the MPC setting. Their work split the operator into two (or even more) different operators. The orders are sent in an encrypted form (secret sharing [7]) to the operators, such that no single operator can recover the encrypted orders. This addresses the problem that we are considering here, as no operator by itself can learn anything about the orders that were submitted. The trust model is depicted in Figure 2.

This trust model is significantly better than the one we have today in which clients must trust the operation, and submit their orders in the clear. Nevertheless, there are few drawbacks with this trust model. First, there is still a possibility of abuse if the operators collude. In such a scenario, the operators can collude and recover the orders by exchanging information. Second, it is unclear who these additional operators are. The work of [3] suggests that the additional operator would be the regular itself. Yet, it is unknown whether the regulators will be able to engage in such a service and becoming an active part in its operation. Finally, the system in [3] leaks the type of the orders, i.e., whether an order is a Buy or Sell, while hiding just the volume and its price. This leakage is significant since the operator(s) can infer imbalances in the orders (e.g. as opposed to Sell orders, Buy orders are less likely to be fully matched) and subsequently imbalances in the market.

Our mechanism. We show a cryptographic protocol for the case of a single operator. Our protocol is based on a variant of FHE, and works, at a high level, as follows: Each client sends its order to the operator in an encrypted form. Using the homomorphic properties of the encryption scheme, the operator can find whether there is a match between the submitted order and previously submitted orders, i.e., the order book. However, it can only obtain this information encrypted. In order to decrypt, we present a sub-protocol for decryption this information (i.e., whether a matching order was found). The operator can then execute the matching orders. Our protocol hides also the type of the order, which requires some subtle treatment in case of partially matched orders.

Disclaimer. This paper was prepared in part for information purposes by the Artificial Intelligence Research group 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.

REFERENCES

[1] Benjamin Bain. 2018. Wall Street Dark Pools to Come Out of Shadows Thanks to SEC. (2018). https://www.bloomberg.com/news/articles/2018-07-18/wall-street-dark-pools-set-to-come-out-of-shadows-thanks-to-sec.

[2] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. 1988. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*. 1–10.

[3] John Cartlidge, Nigel P. Smart, and Younes Talibi Alaoui. 2019. MPC Joins The Dark Side. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security, AsiaCCS 2019, Auckland, New Zealand, July 09-12, 2019*. 148–159.

[4] David Chaum, Claude Crépeau, and Ivan Damgård. 1988. Multiparty Unconditionally Secure Protocols (Extended Abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*. 11–19.

[5] Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*. 169–178.

[6] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*. 218–229.

[7] Adi Shamir. 1979. How to Share a Secret. *Commun. ACM* 22, 11 (1979), 612–613.

[8] Andrew Chi-Chih Yao. 1982. Protocols for Secure Computations (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*. 160–164.

[9] Andrew Chi-Chih Yao. 1986. How to Generate and Exchange Secrets (Extended Abstract). In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*. 162–167.