

FairTraDEX: A Decentralised Exchange Preventing Value Extraction

Conor McMenamin¹, Vanesa Daza², and Matthias Fitzi³

^{1,2}Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

¹NOKIA Bell Labs, Nozay, France

²CYBERCAT - Center for Cybersecurity Research of Catalonia

³IOHK

February 10, 2022

Abstract

An idealised decentralised exchange (DEX) provides a medium in which players wishing to exchange one token for another can interact with other such players and liquidity providers at a price which reflects the true exchange rate, without the need for a trusted third-party. Unfortunately, extractable value is an inherent flaw in existing blockchain-based DEX implementations. This extractable value takes the form of monetizable opportunities that allow blockchain participants to extract money from a DEX without adding demand or liquidity to the DEX, the two functions for which DEXs are intended. This money is taken directly from the intended DEX participants. As a result, the cost of participation in existing DEXs is much larger than the upfront fees required to post a transaction on a blockchain and/or into a smart contract.

We present FairTraDEX, a decentralised variant of a frequent batch auction (FBA), a DEX protocol which provides formal game-theoretic guarantees against extractable value. FBAs when run by a trusted third-party provide unique game-theoretic optimal strategies which ensure players are shown prices equal to the liquidity provider’s fair price, excluding explicit, pre-determined fees. FairTraDEX replicates the key features of an FBA that provide these game-theoretic guarantees using a combination of set-membership in zero-knowledge protocols and an escrow-enforced commit-reveal protocol. We extend the results of FBAs to handle monopolistic and/or malicious liquidity providers, and provide a detailed pseudo-code implementation of FairTraDEX based on existing mainstream blockchain protocols.

1 Introduction

One of the most prominent and widely-used classes of protocols being run on smart-contract enabled blockchains are decentralised exchange (DEX) protocols. DEX protocols allow a specific set of players, whom we call *clients*, to exchange one token for another in the presence of *market-makers* (MMs), who provide liquidity to clients, usually in exchange for a fee. Interacting with a blockchain-based DEX requires a client or MM to first interact with the players who add transactions to the blockchain, known as miners or block producers. These interactions typically reveal a player’s intention to trade to the block producer before the transaction is confirmed on the blockchain, and in doing so present the block producer with what has become known as a miner-extractable value (MEV) opportunity. MEV, first coined in [15], refers to any

email: conor.mcmenamin@upf.edu — vanesa.daza@upf.edu — matthias.fitzi@iohk.io



This Technical Report is part of a project that has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement number 814284

expected profits the miner of a block can extract from other players trying to interact through the miner with the blockchain. This extraction is performed by manipulating the ordering of, injecting, or censoring transactions in prospective blocks.

However, MEV is merely the tip of the iceberg. In blockchain-based protocols such as DEXs, it is also possible for one player to extract value from another by disobeying the protocol instructions, an attack we refer to as *selective participation*. Many decentralised protocols do not bind players to correctly performing protocol actions. In such protocols, players are given a form of *optionality*, the option to correctly perform protocol actions or not, which is typically not intended by the protocol. Optionality is well understood in derivative markets, and is something which must be charged for. Otherwise, players receive this optionality/value at a discount/for free. Given optionality exists in a protocol, this value can be used by one player, and taken from others. As such, value can theoretically be extracted by every player in the blockchain ecosystem, not just the miners. We refer to any value which can be extracted from players in a blockchain-based protocol other than fixed upfront fees as *expected extractable value* (EEV) in line with recent work [23] on the generalisation of extractable value.

A significant advancement in DEX protocols was the advent of *automated market makers* (AMMs) such as Uniswap¹, based on the principle of maintaining a constant ratio of two tokens in a pool which forms the basis for the price of swapping one token for another. AMM pools are built to reflect the market-implied fair price of one token expressed in another, and to provide a permanent source of liquidity for the token swap, accessible by any player in the blockchain ecosystem. Over large time horizons, both of these purposes have been fulfilled. However, on a transaction-by-transaction level view, projects like Flashbots (a direct spin-off to [15]) have identified that AMMs are the main source of known EEV (> 90% of the \$500 million in EEV identified by Flashbots since August, 2020)². Furthermore, Flashbots only observes basic forms of EEV, meaning in reality (and as stated by the Flashbots team³), this amount of EEV is a lower bound for the total amount of value being extracted from clients and MMs alike through participation in AMMs. Although many attempts have been taken to address this significant source of EEV [1, 3, 12, 17, 24], no satisfactory solution has been found. The protocols presented in these works remain vulnerable to basic EEV attacks in the case where all transactions are eventually added to the blockchain (a property we call censorship-resistant), with such attacks outlined in Section 2.

Therefore, there is a clear gap, both in literature and in practice, to provide a DEX protocol which definitively eliminates all sources of EEV. In this paper, we provide such a protocol. We outline FairTraDEX⁴, a DEX protocol based on an existing auction process called *frequent batch auctions* (FBAs) [8] and zero-knowledge (ZK) tools for set-membership. FairTraDEX is, to the best of our knowledge, the first DEX protocol outlining specific practical conditions under which EEV is prevented in a censorship-resistant blockchain. We formally prove this security against EEV by utilising results from ZK literature and game-theory.

Specifically, we show that given a sufficient number⁵ of registered participants in FairTraDEX, the strategy of players trading tokens with each other at the market-implied fair price, minus fees, forms a strict Nash equilibrium. This is compared to DEX protocols in which EEV opportunities exist, where for example, a potential value-extractor responds differently to buy orders vs. sell orders, or once-off clients vs. professional MMs. FairTraDEX replicates the private submission of orders to a trusted third-party in an FBA, revealing no public information to the blockchain protocol other than set of all registered players, and the number of players in an auction. Many protocols in-use today⁶ replicate this exact proving of set-membership in zero-knowledge functionality. These, coupled with our precise pseudo-code implementations (Appendix F.5) for the settling of orders and the utilisation of generic ZK set-membership protocols within FairTraDEX provide a clear guide for the adoption of FairTraDEX on every major smart-contract enabled blockchain.

¹<https://uniswap.org/>

²<https://explore.flashbots.net/> Accessed: 28/01/2022

³<https://explore.flashbots.net/faq> “Are you quantifying all MEV in this dashboard” Accessed: 28/01/2022

⁴A **F**air **T**rade-price **D**ecentralised **E**Xchange.

⁵This number of participants is described in Section 5.3, and used analogously to Zerocash[4], and it’s derivatives. An increase in minted tokens (registrations in FairTraDEX) decreasing the probability that a randomly selected token corresponds to the player who minted it.

⁶ZCash [28], TornadoCash <https://tornado.cash/>, and Semaphore <https://semaphore.appliedzkp.org/>, to name but a few.

1.1 Our Contribution

We first describe a width-sensitive frequent batch auction (WSFBA), an idealised commit-reveal token-exchange protocol between clients and MMs, based on an FBA [8]. In the first phase, clients and MMs anonymously generate orders (price, direction, and an amount of tokens to trade) and markets (bid price, offer price, and token amounts to trade at both prices) respectively. Any client wishing to trade commits to their order, and the maximum width (distance between the bid and ask) of a MM-provided market on which the client will trade. Clients committing to an order with maximum width w only trade if a MM commits a to market with width less than or equal to w . Any player wishing to participate as a MM commits to a market (bid and offer) in a specified minimum size used to ensure the MM cannot perform value extraction attacks (discussed in the proof of Theorem 4.4). In the final phase, all private information regarding client orders and MM markets is revealed. The revealed orders and markets are then settled at a *clearing price CP* which maximises the total (notional) value (described specifically in Section 3) traded by all players. We then prove that the optimal strategy for rational clients is to submit market orders (buy/sell at any price) and for rational non-colluding MMs to show markets with bid, equal to the market-implied fair-value (MIFP) for the swap, equal to the offer. Furthermore, given the decentralised nature of the protocol, we also consider the case where there is a single monopolistic MM. For such a MM, we show the strategy of showing a market geometrically centred around the MIFP with width equivalent to at most the fee clients are willing to pay to participate in the protocol forms a strict Nash equilibrium.

We describe FairTraDEX, a blockchain-based implementation of a WSFBA. To ensure the correct revelation of information as in the second phase of WSFBA, clients and MMs who commit to orders in the first phase must deposit an *escrow* which is destroyed if revelation is not performed correctly. A representation of the information and escrow flow in FairTraDEX prior to order settlement is presented in Figure 1. As commits are recorded on-chain (to enforce the corresponding escrow punishment), we utilise ZK set-membership proofs to allow clients to commit to their orders anonymously. As such, in FairTraDEX, every client must initially register to the protocol, and deposit their escrow in the registration. Then, whenever a client wants to commit to an order, the client only has to prove that they are a member of the set of players who registered in the protocol. Given enough registrations, the probability a client’s ZK set-membership proof and committed order relates to the actual order contents approaches 0 (we formalise this notion in Section 5) such that no other player in the system can see the committed order and use it to infer anything about what the order is. To definitively hide a client’s order-information, orders are committed, including the ZK-proof, by using a relayer. A relayer, as discussed in Section 3.3, is a third-party who receives a fee for including relayed transactions in the blockchain.

This combination of tools serves as a subtle, yet crucial, improvement to previous attempts to implement blockchain-based FBAs [13, 17] where game-theoretic guarantees are dependent on the hiding of order information until all markets and orders have been committed. In these previous attempts, on-chain order commitment reveals client/order-specific information which allows observant counterparties to improve their ability to guess client order information. With sufficiently many registered clients, no counterparty using FairTraDEX can learn information that can be used to bias prices against clients with positive expectancy. As such, EEV is effectively prevented in FairTraDEX when enough clients register to the protocol.

We provide a detailed pseudo-code implementation of FairTraDEX, and given the importance of deploying such an EEV-resistant DEX, we discuss the practical considerations that must be taken when implementing FairTraDEX for real-world usage.

1.2 Organisation of the Paper

Section 2 analyses all previous work related to the construction of EEV-proof DEX protocols. Section 3 outlines the financial terminology used in the paper, the player definitions (3.1), the ZK primitives (3.2) and relayer functionalities (3.3) needed to formally reason about FairTraDEX. Section 4 defines the ideal WSFBA functionality, and identifies the strategies of both clients and MMs. Section 5 maps the ideal WSFBA functionality to a series of algorithms which form the FairTraDEX protocol. It is then proved that rational players follow FairTraDEX, and as such, implement a WSFBA. We conclude in Section 6. The Appendices include an extended related work (App. A), extended discussion on ZK literature (App. B), proofs not included in the main body of the paper (App. C), a detailed description of the FairTraDEX algorithms (App. D), a description of a clearing price verifier for settling orders in FairTraDEX (App. E),

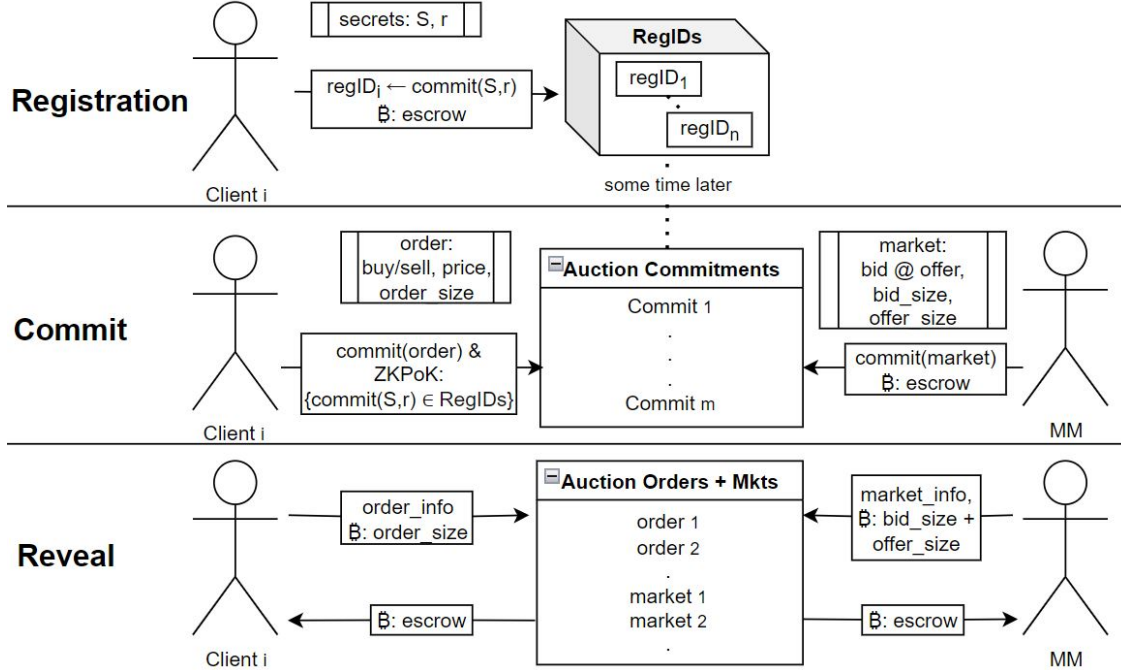


Figure 1: FairTraDEX phases before order settlement. \mathfrak{B} : indicates the transfer of some tokens, but not necessarily in the same denomination.

the pseudo-code for encoding FairTraDEX as a smart-contract (App. F), and notes on FairTraDEX (App. G).

2 Related Work

The main works aimed at protecting DEX users from EEV either focus on preventing front-running of orders [3, 12], the fair ordering of transactions based on their delivery time [24], or on hiding client trade information until the trade has been committed to the blockchain [1, 13, 17]. Of these works, the closest to our proposal are [3, 12, 13, 17]. All of these works critically depend on honesty from MMs, auction operators and/or the block proposers. In Appendix A, we provide high-level explanations of the DEX protocols while in this section, we briefly outline how, when all players in the respective DEX protocols are rational, EEV opportunities exist.

In [12], the MM is always allowed to see orders from clients and can choose to abort them. In Section 2.4 of [12], it is argued that MMs are happy to trade against all orders, including informed orders. Furthermore, it is assumed that clients are independent, with random information. This is not true in real-world trading environments, and as such, Theorem 3 may not hold in practice. The paper references [16] as justification for the quality of price/service provided by the MM, however there is a subtle but crucial difference between the games in [16]: in [16] the client has the final decision on whether or not to execute the trade, while in [12] the MM has the final decision. This optionality has an implicit, but not explicit, cost for the client and provides a source of EEV to the MM. In FairTraDEX, no optionality is given to the MM or clients, while given enough registered clients, the direction of any individual client remains hidden until the trade has been committed. As such, FairTraDEX is able to benefit from the results of [16] in a single MM game, that is, liquid (tradeable and of a client-specified width) markets centred around the pre-trade market-implied fair value (MIFP) when a trade is accepted by the MM.

In P2DEX [3], clients deposit tokens to trade publicly in the same time frame as the order matching takes place, exposing clients to all of the standard identity- and directional-based EEV exploits unless the number of clients per-auction in [3] is equal to the total number of registered clients in FairTraDEX. Separating token deposit and identity revelation from a client’s commitment to a specific auction are important advancements

used in FairTraDEX to protect against EEV. Furthermore, at least one of the servers in charge of fairly executing orders is required to be honest-by-default. If the servers, a minority subset of players in the P2DEX protocol, are rational and monopolised/colluding, the servers can front-run orders. In FairTraDEX, such value-extraction is prevented by keeping all order information hidden until every order in a particular settlement round has been committed. Furthermore, the set of servers act as a point-of-failure as server participation is required to finalise order-settlement. No subset of players in FairTraDEX can prevent the matching of correctly-revealed orders, while in P2DEX, any majority of servers can prevent order-matching.

Similar commit-reveal protocols to FairTraDEX for blockchain-based token-exchange are proposed in [13, 17]. The protocol in [17] is exposed to several game-theoretic exploits which contradict its protection against front-running. These include the necessity to reveal order direction a-priori, and the non-trivial handling of the linkability between commitments and account-balances. The protocol also depends on an operator who does not participate in token-exchange, gains exclusive access to order information, and is depended on for protocol completion. In [13], clients commit their own orders to the blockchain, revealing their identities, and corresponding token balances/execution patterns which can be used by a basic professional MM to skew prices and extract value from the client. Both of [13, 17] aim to protect unmatched orders from being revealed, but to do so, both protocols depend on a third party selected before the auction begins to execute order-matching (no one else in the ecosystem can finalise the auction, a single point of failure). Both protocols assume that the operator does not reveal unexecuted order information.

3 Preliminaries

This section introduces the terminology and definitions necessary to understand the main results of the paper. By $negl()$ we denote any function $f : \mathbb{N} \rightarrow \mathbb{R}$ that decreases faster than any (positive) polynomial p : $\forall p \exists \lambda_0 \in \mathbb{N} : \forall \lambda > \lambda_0 : f(\lambda) < \frac{1}{p(\lambda)}$.

For protocol correctness, we must assume that some of the involved players may be malicious trying to force the protocol into incorrect execution, and without any direct benefit for themselves. However, for the game-theoretic part of the analysis, we assume that all players are rational. Accordingly, the analysis of our protocol is based on two security parameters, a cryptographic security parameter κ bounding the probability that the protocol execution is incorrect by $negl(\psi)(\kappa)$; and a game-theoretic extractable-value parameter ψ bounding the extractable value by any player by $negl(\psi)(\psi)$.

3.1 Definitions and Terminology

In this section, we review the financial and game-theoretical terms used in this paper. Although not mandatory for all readers, this section serves as a useful reference point towards understanding the results and discussions that follow.

- *Decentralized Exchange (DEX)*: A distributed marketplace which allows players to swap one token for another.
- *Limit Order*: Specifies an amount of tokens to be bought (sold), and a maximum (minimum) price at which to buy (sell) these tokens. This price is known as the *limit price*.
- *Market Order*: Specifies an amount of tokens to be sold, but no limit price. Market orders are to be executed immediately at the best available price based on the liquidity of buy orders.
- *Market*: A market in a DEX between two tokens A_{tkn} and B_{tkn} consists of two limit orders, a *bid* and *offer*. When the market is quoted from token A_{tkn} to B_{tkn} , the offer price indicates the quantity of token B_{tkn} a player receives for 1 token A_{tkn} , while the bid price indicates the quantity of token B_{tkn} a player must sell to receive 1 token A_{tkn} . In this paper, we represent such a market as *bid @ offer*, with $0 < bid \leq offer$.
- *Direction*: With respect to an order on a market quoted from token A_{tkn} to B_{tkn} , if the order is trying to buy token B_{tkn} , the direction is *buying*, while if the order is trying to sell token B_{tkn} , the direction is *selling*.

- *Notional Value*: The value of a set of tokens expressed in some common reference token. In this paper, we use the symbol \mathfrak{B} as the reference token in which we measure notional, and with which we reason about utility.
- *Market-Implied Fair Price* (MIFP, denoted y) [8]: The MIFP of a token/token swap is a publicly observable signal (all players in the system observe the MIFP) from all players in the system which is perfectly informative of the fundamental fair price of the underlying token/token swap ⁷. Moreover, a random order of fixed notional generated by a player in the system is equally likely to buy or sell tokens at the MIFP, distributed symmetrically around the MIFP.
- *Reference Price* (y_{ref}): For a market $bid @ offer$, the reference price y_{ref} is the price such that $\frac{bid}{y_{ref}} = \frac{y_{ref}}{offer}$. This y_{ref} is known as the *geometric mean* of bid and $offer$.
- *(Market) Width* (w): For a market $bid @ offer$, the width is calculated as $w = \frac{offer}{bid}$ (As such $w \geq 1$). It can be seen that $\frac{y_{ref}}{\sqrt{w}} = bid @ offer = \sqrt{wy_{ref}}$
- *Multiplicative Market-Impact Coefficient* (δ): If the pre-trade MIFP is y , the expected post-trade MIFP given a buy order is δy for some $\delta \geq 1$, while the expected post-trade MIFP given a sell order is $\frac{y}{\delta}$. It is common in literature to implicitly assume $\delta = 1$, indicating trades have no impact on MIFP. In reality, trades act as signals for the fundamental fair price, and as such, we consider trade impact in our analysis. In this paper, for clients trading on a market of width w with $y_{ref} = y$, we assume $\delta < w$.⁸
- *Forward Price*: This is the price at which a seller delivers a token to the buyer at some predetermined date. In any exchange protocol without instantaneous delivery, the forward price at expected delivery time is the price at which trades should happen. The difference between current (spot) price and forward price is known as *carry*, and can be due to storage/opportunity costs, interest rates, etc. In this paper, we set carry to 0 for complexity and ease-of-notation purposes.
- *Strict Nash Equilibrium* [30]: Consider a set of non-cooperative players P_1, \dots, P_n , with strategies (series' of actions) str_1, \dots, str_n describing the actions which each player takes throughout a particular protocol. These strategies form a strict Nash Equilibrium if any individual player deviation from these strategies strictly reduces that player's utility.
- *Strong Incentive Compatible in Expectation* (SINCE) [26]: With respect to a protocol strategy, a protocol strategy is SINCE for a player if all other strategies have strictly smaller expected utility.

Definition 3.1. A *client* is any player in a DEX protocol who, for an MIFP y , there exists some *minimum client utility* $f_{mcf} > 1$ such that client buyers (sellers) have positive expected utility to trade for or below $\sqrt{f_{mcf}}y$ (at or above $\frac{y}{\sqrt{f_{mcf}}}$). A *market maker* (MM) is a player in a DEX protocol with large supplies of all tokens, who has positive expected utility trading with clients on markets of any width $w > 1$ with reference price equal to the MIFP.

Every player in a DEX can be classified as either a client or MM (sometimes referred to as a liquidity provider), with some players classifiable as both a client and a MM, although they are treated separately in the analysis that follows. In creating a DEX protocol, we must ensure that there is a strict Nash Equilibrium in which rational client and MMs trade at the MIFP with pre-determined fees bounded by the utility gained by any individual player, effectively removing EEV.

As stated in Section 1, FairTraDEX is based on an FBA [8]. We define an FBA here using the terminology of our paper. FBAs are used in many of the largest centralised exchanges ⁹. As FBAs were initially intended for

⁷For a market quoted from token A_{tkn} to token B_{tkn} , given MIFPs with respect to \mathfrak{B} of y_A and y_B respectively, the MIFP of the market from A_{tkn} to B_{tkn} is $y_{A \rightarrow B} = \frac{y_B}{y_A}$. As y_A represents the amount of \mathfrak{B} s equivalent to 1 A_{tkn} , so $y_{A \rightarrow B}$ is the amount of A_{tkn} s per 1 B_{tkn} . For n B_{tkn} s, given an MIFP of $y_{A \rightarrow B}$, this is equivalent to $ny_{A \rightarrow B}$ A_{tkn} s.

⁸Without this assumption, the proof of Theorem 4.4 will not hold in general. We leave deeper analysis of δ as future work.

⁹FCA <https://www.fca.org.uk/publications/research/periodic-auctions>, CBOE https://www.cboe.com/europe/equities/trading/periodic_auctions_book/, ESMA https://www.esma.europa.eu/sites/default/files/library/esma70-156-1035_final_report_call_for_evidence_periodic_auctions.pdf

a centralised setting, we consider them being run by a trusted third party (TTP) who enforces the correct participation of all players. In FairTraDEX, the key TTP functionalities needed to instantiate an FBA are replicated using ZK set-membership proofs, incentivisation and a blockchain as a censorship-resistant bulletin-board.

Definition 3.2. A *frequent batch auction* (FBA) (sometimes referred to as a periodic auction) involves clients and MMs privately submitting either limit or market orders to the TTP. These orders are collected until a specified deadline. After this deadline, the orders are settled at the clearing price. A single *clearing price* is chosen which maximises the total notional traded based on the specified sizes and prices of all orders. If there is more supply (quantity of tokens being sold) at the clearing price than demand (quantity of tokens being bought), all sell orders offered at the highest price at or below the clearing price are pro-rated based on size such that supply equals demand at the clearing price. Similarly, if there is more demand than supply at the clearing price, all buy orders bid at the lowest price at or above the clearing price are pro-rated based on size such that demand equals supply at the clearing price. Any limit buy orders below/sell orders above the clearing price are not executed.

There are two key differences between this definition and the specification in [8]:

1. In our definition, if an order is not fulfilled, it is revealed with any tokens not being sold returned to the seller. This does not affect the game-theoretic guarantees of the paper, as the results in [8] only depend on the hiding of order information while players are submitting orders to the auction.
2. As every order in our auction must be submitted independently for each auction, there is no time priority applied when pro-rating orders in case of a supply-demand imbalance. This is a sub-case of the FBAs as defined in [8], and consequently, our protocol retains the same game-theoretic guarantees.

3.2 Zero-Knowledge Primitives

The aim of this section is to outline the generic *non-interactive zero-knowledge* (NIZK) tools for set membership as used in this paper, such as those stemming from papers like [4, 6, 21, 22, 28]. We generalise these formal works, allowing for the adoption of any secure NIZK set-membership protocol into FairTraDEX, as we only require a common functionality that is shared by all of them. Further elaboration on these protocols is deferred to Appendix B.

In a nutshell, set membership is the problem of a prover P proving to a verifier V (or a set of verifiers) that an element x is in a (usually public) set S . The prover computes a proof π and sends it to the verifier, who should be able to efficiently check its validity. Due to the large number of applications in the last decade, especially related with blockchain technology, research has intensively focused on two fronts: reducing the amount of resources that are needed to check membership (scalability), and guaranteeing membership while keeping x secret (privacy).

The zero-knowledge proofs used in FairTraDEX allow, for a given set of commitments Com to user-generated secrets, that any user knowing the secret corresponding to a commitment $com \in Com$ can prove the knowledge of a secret corresponding to a commitment in the set, without revealing which secret, or commitment. Moreover, we require that more than one proof relating to the same commitment is identifiable by a verifier.

To participate in FairTraDEX, clients privately generate two bit strings, the *serial number* S and *randomness* r , with $S, r \in \{0, 1\}^{O(\kappa)}$. To describe FairTraDEX we define a commitment scheme f_{com} , a set membership proof scheme $SetMembership$, a NIZK proof of knowledge scheme $NIZKPoK$ and a NIZK signature of knowledge scheme ($NIZKSoK$). We do not specify which instantiation of these schemes to use, as the exact choice will depend on several factors, such as efficiency, resource limitations and/or the strength of the assumptions used.

- $f_{com}(m)$: A deterministic, collision-resistant function taking as input an arbitrary length string $m \in \{0, 1\}^*$, and outputting a string $com \in \{0, 1\}^{O(\kappa)}$.
- $SetMembership(com, Com)$: Compresses a set of commitments Com and generates a membership proof that com is in Com if and only if $com \in Com$.

- $NIZKPoK(r, S, Com)$: For a set of commitments Com , returns a string S and NIZK proof of knowledge if and only if the person running $NIZKPoK()$ knows a r producing a proof when running $SetMembership(f_{com}(S||r), Com)$. In FairTraDEX, this revelation identifies to a verifier when a proof has previously been provided for a particular, albeit unknown, commitment^{10 11} as the prover must reproduce S . This is used in FairTraDEX, in conjunction with an escrow, to enforce the correct participation of both clients and MMs.
- $NIZKSoK(m)$: Returns a signature of knowledge that the person who chose m can also produce $NIZKPoK$.

3.3 Relayers

A fundamental requirement for transaction submission in blockchains is the payment of some transaction fee to simultaneously incentivise block producers to include the transaction, and to prevent denial-of-service/spamming attacks. However, in both the UTXO- and account-based models, this allows for the linking of player transactions, balances, and their associated transaction patterns. With respect to DEX protocols, if clients are required to deposit money into a UTXO/account before initiating a trade, any other player in the system can infer who the client is, what balances the client owns, what transactions the client usually performs, etc., and use this information to give the client a worse price.

To counteract this, we utilise the concept of *transaction relayers*¹². In the smart-contract encoding of FairTraDEX (App. F.5), clients must publicly register to a smart contract, and in doing so, deposit some escrow. In addition to this escrow, we also require the clients to deposit a relayer fee. When the client wishes to submit a transaction anonymously to the blockchain, the client publishes a proof of membership in the set of registered clients to the relayer mempool, as well as the desired transaction and a signature of knowledge cryptographically binding the membership proof to the transaction, preventing tampering. As the relayer can verify the proof of membership, the relayer can also be sure that if the transaction is sent to the FairTraDEX contract, the relayer will receive the corresponding fee. With this in mind, a relayer observing the client transaction includes it in a normal blockchain transaction, with the first relayer to include the transaction receiving the fee. As such, relayers are a straightforward extension of the standard transaction-submission model. Furthermore, if the proof of membership is NIZK and the message is broadcast anonymously (using the onion routing (Tor) protocol¹³ for example), the relayer can only infer that the player sending the transaction is a member of the set of clients.

4 Frequent Batch Auctions

In this section we outline the properties of an idealised FBA towards constructing a variation of FBAs which we define as a width-sensitive FBA. Width-sensitive FBAs maintain the desirable properties of FBAs with respect to optimal strategies for MMs and clients, while also adding important protections for clients in a decentralised setting where monopolistic MMs may exist.

4.1 Game-Theoretic Guarantees of a Frequent Batch Auction

In this section we investigate the properties of an FBA between rational MMs and rational clients, where MMs do not know the desired trade direction of the clients.

We first restate, using the terminology from this paper, the main result from [8] which applies to our game-theoretically equivalent definition of an FBA. To do this, we let D represent the net trade imbalance of clients in a particular instance of an FBA in terms of \mathfrak{B} . A positive D indicates a client buy imbalance (more client buyers than sellers of the swap), while a negative D indicates a client sell imbalance. We require a finite bound on the absolute imbalance, which we denote $Q_{imb} < \infty$, for the existence of optimal MM

¹⁰In Z(ero)cash [4], serial numbers are also used for the same purpose; to prevent double-spending of coins.

¹¹In Semaphore [22], a similar technique using one-time nullifiers prevents the reuse of a commitment for protocols where one interaction per commitment is required, such as one-person one-vote voting.

¹²Ox <https://0x.org/docs/guides/v3-specification>, Open Gas Station Network <https://docs.opengsn.org/>, Rockside <https://rockside.io/>, Biconomy <https://www.biconomy.io/>

¹³<https://www.torproject.org/>

strategies. As in [8], we assume that $|D| \leq Q_{imb}$, and in-keeping with the notion of an MIFP, D is symmetric around 0 at the MIFP.

Theorem 4.1. [8] For an FBA with at least two non-cooperative MMs, there is a strict Nash equilibrium where clients only submit market orders and MMs show a market of width 1 (*bid = offer*) centred around the MIFP in size greater than Q_{imb} .

This is a useful result in the case of at least two non-cooperative MMs, with clients receiving a game-theoretic guarantee that they can exchange one token for another at the MIFP in expectancy in an FBA. Furthermore, as MM liquidity is greater than the net client trade size, the implicit impact to these trades in [8] is bounded by the width, which is 1. As clients have a strictly positive utility for exchanging tokens, this is equivalent to clients always having positive expectancy to participate in an FBA. However, it is also shown in this equilibrium that MMs have 0 expected utility. A basic adjustment to the protocol in that setting would then be to charge clients a fee for the service and pro-rate these fees to the MMs to ensure the long-term participation of MMs.

4.2 Width-Sensitive Frequent Batch Auctions

The important assumption in the above result is the presence of at least two non-cooperative MMs. In a decentralised setting, this can be seen as insufficient. To reach a similar equilibrium in the presence of a monopolistic MM, we must amend the basic FBA protocol. In this section, we define a *width-sensitive* FBA (WSFBA) to handle monopolistic MMs, while retaining the desirable properties of an FBA in the presence of two or more non-cooperative MMs. Towards this definition, we now describe how trades at prices other than the MIFP affect the MIFP, and (re)introduce the variables necessary to define a WSFBA.

When considering a trade against an order with unknown direction, we assume a multiplicative market-impact coefficient of δ . Furthermore, in the results that follow, for a token swap between any two tokens A_{tkn} and B_{tkn} with pre-trade MIFP $y_{A \rightarrow B} = \frac{y_B}{y_A}$, and symmetric impact coefficient δ , the post-token swap trade MIFP for B_{tkn} has multiplicative market-impact coefficient of $\sqrt{\delta}$, while the post-token swap-trade MIFP for A_{tkn} has multiplicative market impact coefficient of $\frac{1}{\sqrt{\delta}}$. Specifically, if a token swap with pre-trade MIFP $y_{A \rightarrow B}$ has an expected post-token swap trade MIFP in the case of a client buyer of $\delta y_{A \rightarrow B}$, given a client buyer of the swap from A_{tkn} and B_{tkn} , the expected post-token swap trade MIFP of A_{tkn} is $\frac{1}{\sqrt{\delta}} y_A$ while the expected post-token swap trade MIFP of B_{tkn} is $\sqrt{\delta} y_B$.

In the presence of a single rational MM, we need to utilise the value gained by clients for exchanging token. That is, recall from Section 3, clients in our protocol observe a positive utility of at least the minimum client fee f_{mcf} for exchanging tokens. We introduce a notional amount Q_{not} which upper-bounds the total notional allowable in the auction through limit and market orders (as total auction notional must be greater than auction imbalance, we have $Q_{not} \geq D$), and which lower-bounds the notional at the bid and offer of a market. This upper-bound on the notional of orders allowable in a particular WSFBA being the lower-bound on the notional of a MM's bid and offer is used when reasoning about the strategy of a monopolistic MM in Theorem 4.4. Informally, it ensures a monopolistic MM providing a market cannot synthetically remove the market bid or offer from the auction (through self-trading) without preventing clients interacting with the MM (a strictly dominated strategy). To self-trade all of either the market bid or offer and change the MM's net order would use all of the available market notional bandwidth. We now define a WSFBA.

Definition 4.2. A *width-sensitive frequent batch auction* (WSFBA) with notional bound Q_{not} involves MMs submitting markets to the TTP with total notional on the bid and offer of at least Q_{not} . Clients and MMs also privately submit limit and market orders to the TTP including a requested maximum width from the tightest MM, above which the order is not executed. This can be set to *any*, indicating the client will trade regardless of width. Orders are collected until a specified deadline while the total notional of orders is less than or equal to Q_{not} . After this deadline, clients orders with requested width greater than or equal to the tightest MM width, along with a randomly-selected market from the tightest provided markets, are settled at the clearing price. A single clearing price is chosen which maximises the total notional traded, and then minimises the net trade imbalance, based on the specified sizes and prices of all orders. If there is more supply at the clearing price than demand, sell orders at the highest price at or below the clearing price are pro-rated based on size such that supply equals demand at the clearing price. Similarly, if there is more

demand than supply at the clearing price, buy orders at the lowest price at or above the clearing price are pro-rated based on size such that demand equals supply at the clearing price. Any limit buy orders below/sell orders above the clearing price are not executed.

The key differences between a conventional FBA and a WSFBA are the specification of MM widths by clients, the minimum MM notional requirement on the bid and offer which also stands as the maximum allowable size of the auction, the exclusive access of the MM with the tightest market width to the auction, and the requirement for the clearing price to minimise the imbalance over all prices which maximise the notional traded. Minimising imbalance is a small optimisation which produces a reasonable and precise clearing price when MMs do not show width 1 markets as in an FBA. The other amendments are intended to protect clients against monopolistic MMs, and are discussed in the proceeding section.

4.3 Properties of Width-Sensitive Frequent Batch Auctions

In Theorem 4.4, we show that the strict Nash equilibrium in a WSFBA under more than 1 non-cooperative MMs is equivalent to that of an FBA. The case of a single monopolistic MM is more complex. First, we observe that an MM in a WSFBA always shows a market with reference price equal to the MIFP. In the proceeding lemmas, proofs omitted from the main body of the paper are included in Appendix C.

Lemma 4.3. For an MM in a WSFBA between A_{tkn} and B_{tkn} with MIFP equal to $y_{A \rightarrow B} = \frac{y_B}{y_A}$ and a client order of notional $X\mathfrak{B} > 0$, it is strong incentive compatible in expectation to show a market with reference price $y_{ref} = y_{A \rightarrow B}$ for any fixed width $w \geq 1$.

This result is independent of the choice of width and impact coefficient. However, it assumes that the MM trades with the client on either the bid or the offer. With respect to a WSFBA without notional restrictions and a monopolistic MM, if clients submit market orders, there are fringe cases (large imbalances) which incentivise MMs to show markets far from the MIFP. Without restrictions on the number of markets allowed into the auction or on notional of client orders, a MM can manipulate the width restriction and increase the expected utility of the MM. Removing these restrictions from a WSFBA makes for interesting future work.

Recall, clients have a strictly positive utility to exchange tokens described by the minimum client fee f_{mcf} , which is equivalent to being strongly incentivised to trade on a market with reference price y_{ref} and width $w \leq f_{mcf}$. With this in mind, we can now apply the main result of [8] to a WSFBA.

Theorem 4.4. For a WSFBA, the strict Nash equilibria strategies given the number of non-cooperative MMs submitting markets being N are:

- $N = 1$: Clients submit market orders of requested width f_{mcf} and the MM shows a market of width at most f_{mcf} with reference price equal to the MIFP in size of at least Q_{not} .
- $N \geq 2$: Clients submit market orders of requested width greater than 1 and MMs show a market of width 1 with reference price equal to the MIFP in cumulative size of at least Q_{not} .

Proof. We now investigate each of the cases described in the theorem statement in terms of the number of non-cooperative MMs N .

$N = 1$: Consider first the strategy of a client. For buy orders, the strategy of submitting a limit order with price p less than $\sqrt{f_{mcf}}y$ is dominated by all prices greater than p and less than or equal to $\sqrt{f_{mcf}}y$. For limit sell orders, this limit is $\frac{y}{\sqrt{f_{mcf}}}$. As such, the equilibrium for clients involves submitting orders equivalent to a market of width equivalent to at least f_{mcf} with reference price equal to the MIFP. If a client knows a MM submits a market of width less than or equal to f_{mcf} with reference price equal to the MIFP, this strategy is further dominated by submitting a market order with requested width f_{mcf} , as market orders strictly increase the client's probability of trading. Furthermore, any strategy for a client which involves trading on a price outside $[\frac{y}{\sqrt{f_{mcf}}}, \sqrt{f_{mcf}}y]$ is strictly dominated by not trading. As such, the only possibilities for equilibria can occur on a market of $\frac{y}{\sqrt{f_{mcf}}}$ @ $\sqrt{f_{mcf}}y$. Furthermore, the submission of market orders (increasing probability of trading) with requested width f_{mcf} is strictly dominant if the MM shows a market with reference price equal to the MIFP.

Considering now the MM strategy, we can immediately see the MM only submits one market, as only the tightest market in every auction is included for settlement. For the MM submitting a market $market = bid@offer$, observe that the MM does not submit an order which bids higher than or equal to $offer$, or offers lower than or equal to bid . By definition, the total notional of orders in a WSFBA is bounded by Q_{not} . For the MM to synthetically offer through bid for example, the MM needs to first trade with all notional bid at bid , which is lower-bounded by Q_{not} . Therefore, if the MM wants to trade on a price other than those indicated by $market$, the MM can only do so on prices within $(bid, offer)$. In other words, the MM can only trade on another price by implicitly providing a tighter market within the original market.

We have seen in Lemma 4.3 that a MM trading on a market against a random client shows a market with reference price equal to the MIFP. Furthermore, from Equation 3 we can see that the expected utility of a MM is strictly increasing in width. Any strategy involving a market with width greater than f_{mcf} has 0 expectancy as clients do not trade on such markets, as argued above. Therefore, the MM maximises her expectancy against a random client by showing a market of width f_{mcf} with reference price equal to the MIFP. Against multiple clients, a positive notional imbalance at the MIFP is decreasing in price (resp. a negative notional imbalance is increasing as price decreases), which may cause the MM to provide a market of width less than f_{mcf} .

Given the MM submits a market of width less than or equal to f_{mcf} with reference price equal to the MIFP, this implies clients submit market orders with requested width f_{mcf} , which is a client's strictly dominant strategy. Consider the strategy of a MM providing a market of width less than or equal to f_{mcf} with reference price equal to the MIFP, and the strategy of clients submitting market orders with requested width f_{mcf} . We have shown that any player deviation strictly decreases that player's expectancy, making this a strict Nash Equilibrium.

$N \geq 2$: As MMs in a standard FBA provide markets of width 1 when the width is not a restriction, applying the requested width adjustments of a WSFBA, further incentivising tighter markets, does not change the unique equilibrium of Theorem 4.1. Similarly, as there is a unique clearing price when a width-1 market is submitted, it must also minimise the imbalance over prices that maximise total notional traded. The restriction on the notional of markets in a WSFBA is in line with the inequality of Theorem 4.1. Furthermore, any requested width > 1 in this equilibrium ensures a client's order trading through the MIFP is included in the final auction settlement, with the maximum allocation occurring when a client submits a market order. \square

Theorem 4.4 identifies that clients always submit market orders, and in settings where it is unclear whether there is a single monopolistic MM, or many non-cooperative MMs, it can be seen that clients always submit market orders with requested width f_{mcf} .

5 FairTraDEX

In this section we construct the FairTraDEX protocol as a sequence of algorithms. We then provide a series of results regarding the incentive compatibility of these algorithms with the goal of proving FairTraDEX instantiates a WSFBA, and that following the protocol is SINCE.

In Section 4 we constructed a WSFBA using a TTP to enforce correct player balances, order sizes, revelation of orders, correct calculation of the clearing price and the settlement of orders. In a decentralised setting with rational players, such a TTP does not exist. However, we do have access to censorship-resistant public bulletin boards in the form of blockchain-protocols. As discussed in the Section 1, these bulletin boards have many caveats such as the ordering of transactions based on transaction send time not being preserved (transaction re-ordering attacks). However, if we are able to bound the delay of updates being added to such a bulletin board (transactions being confirmed on the blockchain), we can implement a WSFBA in such a setting. We now formalise the model we use when describing FairTraDEX.

5.1 System Model

1. All players P_1, \dots, P_n are members of a blockchain-based distributed ledger, and a corresponding PKI.

2. The ledger is represented by a linear blockchain with its state progressing by having new blocks sequentially appended. For simplicity, we assume instant finality of blocks meaning that such an appended (valid) block cannot be replaced at any later point in time.¹⁴
3. A transaction submitted by a player for addition to the blockchain (either directly or relayed) while observing blockchain height H , is included (and thus finalised) in a block of height at most $H + T$, for some known $T > 0$, given that the transaction remains valid for sufficiently many intermediate ledger states¹⁵.
4. The public NIZK parameters are set-up in a trusted manner¹⁶.

We do not make any assumptions regarding transaction ordering in blocks. Specifically, the order in which transactions are executed is at the discretion of the block proposer.

A further consideration when choosing the value of T needs to be the cost to censor a transaction for T blocks vs. the maximum loss incurrable by either a client or MM by allowing a trade to be executed. In a WSFBA, notice that this max loss is Q_{not} for all players. For a transaction tx with fee f_{tx} (not to be confused with the minimum client fee) payable to the producer of the block in which tx is contained, and producers not participating as MMs/clients, a rational player can only bribe block producers to censor a transaction paying f_{tx} for $T \leq \frac{Q_{not}}{f_{tx}}$ blocks. This is because a producer needs to receive at least f_{tx} to censor tx .

If block producers are participating as MMs/clients, we need to adjust T . Let $0 < \alpha < 1$ bound the fraction of blocks produced over chains of length greater than T by a MM responding to/the set of clients requesting trades in a particular instance of a FBA (we need to consider all clients in a request phase, as they may all have the same direction, and as such, some positive expectancy to preventing a MM revelation). We need to increase T by a factor of $\frac{1}{1-\alpha}$ (similar to the methodology behind the Chain Quality property in [18, 31]).

5.2 FairTraDEX Algorithms

Each player P_i owns (has exclusive access to) a set of token balances bal_i which are stored as a global variable. For a token tkn , $bal_i(tkn)$ is the amount of token tkn that P_i owns. Keeping the notation from Section 3.2, outputs included in round brackets $()$ are known only to the player running the algorithm, with all other outputs posted to the public bulletin board, updating existing variables/balances where appropriate. Algorithm outputs are not signed, so players observing the output of an algorithm instance can only infer information about the player running the algorithm from public outputs and any corresponding global variable updates.

We now outline FairTraDEX as a set of algorithms: `Setup()`, `Register()`, `CommitClient()`, `CommitMM()`, `RevealClient()`, `RevealMM()` and `Resolution()`. A FairTraDEX instance is initialised by running `Setup()`, and proceeds indefinitely in rounds of three distinct, consecutive phases: *Commit*, *Reveal* and *Resolution*, each of length T blocks (see Section 5.1). We provide here the intuition to the algorithms of FairTraDEX, with a detailed explanation of each algorithm provided in Appendix D.

Players in the underlying blockchain protocol can enter FairTraDEX as clients by running an instance of `Register()`, which for a given client deposits an escrow $escrow_{client}$, and generates private information $(S, r \in \{0, 1\}^{O(\kappa)})$ which is used in `CommitClient()` to prove that the client indeed deposited an escrow, without revealing which deposit.

In the Commit phase, all players can run any number of `CommitClient()` and/or `CommitMM()` instances. `CommitClient()` generates a client order, commits to that order publicly and proves in ZK that the player deposited an escrow. If such a proof cannot be generated, or a proof has already been generated for the same S , no order can be committed. A correctly run `CommitMM()` instance generates a market for a prospective MM, commits to that market publicly and deposits an escrow $escrow_{MM}$.

In the Reveal phase, players can run any number of `RevealClient()` and/or `RevealMM()` instances. `RevealClient()` publishes an order generated through `CommitClient()`, returning the escrow corresponding to

¹⁴An example for instant finality is Algorand [11] which stands in contrast to, e.g., Bitcoin which only guarantees eventual finality.

¹⁵This can be seen as a ‘block-based’ variant of the time-based *liveness* property defined in [18, 31].

¹⁶Such as a Perpetual Powers of Tau ceremony, as used in Zcash <https://zkproof.org/2021/06/30/setup-ceremonies/>

the `CommitClient()` instance, and as such the `Register()` instance, to the client. `RevealMM()` publishes a market corresponding to a `CommitMM()` instance, and returns the corresponding escrow. Both Reveal phase algorithms assert that the client and MM have sufficient token balances to submit their order and market respectively. These assertions are also ensured in the Commit phase, but must be rechecked to ensure correct balances at the point of token transfer.

In the Resolution phase, any number of `Resolution()` instances can be run. The first correct `Resolution()` instance selects the tightest market from the set of revealed markets, *revealedMkts*, for inclusion in order settlement, and any tie-breaks settled using $f_{com}(revealedMkts)$, as a random seed¹⁷. The clearing price which maximises notional traded, and then minimises the notional imbalance of the remaining market and orders is computed¹⁸. Orders and markets are then settled based on this clearing price. Finally, the arrays tracking active commitments, orders and markets *clientCommits*, *MMCommits*, *revealedOrders*, *revealedMkts* are cleared, so unsuccessfully revealed commitments during this round cannot be used to run `RevealClient()` or `RevealMM()` in future rounds. This effectively destroys the deposited escrows of such commitments. .

5.3 Properties of FairTraDEX

We now argue that FairTraDEX possesses all of the necessary properties to instantiate a WSFBA, and discuss the practical implications of these properties. As the `Register()` and `CommitClient()` algorithms are constructed analogously to the `Mint()` and `Spend()` functions in [28], we can make use of the results as provided therein. These can be stated informally as:

1. **Linkability:** Consider a player P_j , a set of registrations *RegIDs* to which P_j does not know the privately committed values, and a valid ZK signature of knowledge π and serial number S corresponding to some $regID_i \in RegIDs$. P_j in has no advantage in linking π and S to the corresponding $regID_i$ over probability $\frac{1}{|RegIDs|} + negl(\kappa)$.
2. **Double-spending:** Given a set of registrations *RegIDs*, and any number of valid (π, S) pairs corresponding to elements in *RegIDs*, it is computationally infeasible to generate a new serial number S' and corresponding valid proof of registration π' in *RegIDs*.

Given that all players in the system are registered as clients, by definition of MIFP, the expected trade imbalance implied by their orders is 0. However, in reality, we cannot expect this ‘full level’ of client participation¹⁹, meaning less client registrations typically results in a greater advantage for rational players in predicting the implied trade imbalance of committed orders.

To account for this in our analysis, we thus introduce n_ψ denoting the minimal number of registrations required to guarantee that EEV is $negl(\psi)(\psi)$. Note that in certain blockchain systems, as the total number of players may be unknown to players within the system, precisely defining n_ψ may not be possible. In that sense, our analysis demonstrates achievability under a sufficient level of registration, but not necessarily that a client can detect whether this level is met in a given auction instance.

In practice, a client’s decision whether or not to commit to an order in FairTraDEX will be based on heuristics involving the number n_c of observed client registrations, noting that non-negligible EEV may be tolerable if the total expected participation fees are less than f_{mcf} .

Observation 5.1. It can be seen that FairTraDEX with at least n_ψ previous `Register()` calls implements a WSFBA when all players follow the protocol. In the Commit phase, `CommitClient()` specifies a client

¹⁷Given all markets are revealed, the final value of *revealedMkts*, and as such $f_{com}(revealedMkts)$, is unpredictable in the presence of two or more non-cooperative MMs. We prove in Lemma 5.4 that all MMs running `CommitMM()` also run `RevealMM()`. The blockchain-based implementation of this function is described in Appendix F.3

¹⁸A precise algorithm for verifying the clearing price is included in Appendix F.5, Algorithm 4, and described in Appendix E. The intuition behind the algorithm is as follows: Given more tokens are sold than bought at the proposed price, it can be seen that checking the next price point lower, first for higher traded notional, and then for a greater or equal absolute imbalance is sufficient to verify the proposed price is a valid clearing price. The equivalent check at the next price point above holds when more tokens are bought than sold at the proposed price.

¹⁹Practical approaches to ensure the equivalent to n_ψ `Register()` calls are taken in existing anonymity protocols. For example, Tornado Cash rewards players proportionally to the (Tornado Cash equivalent of the) number of `Register()` calls, as well as the length of time between calling `Register()` and `CommitClient()`. <https://torn.community/t/anonymity-mining-technical-overview/15> Accessed: 07/02/2022

order which is committed to, while `CommitMM()` specifies a market which is also committed to. As clients commit to these orders and sign this commitment using a NIZKSoK, nothing is revealed about the client’s order, as there are at least n_ψ `Register()` calls. This is equivalent to privately submitting the order.

`CommitMM()` and `RevealMM()` ensure MMs provide the equivalent of at least Q_{not} notional on the bid and offer. Furthermore, as MMs are indistinguishable under Definition 3.1, a MM commitment reveals nothing about the liquidity on the bid or offer²⁰.

When the Commit phase is finished, no further orders or markets can be submitted for that auction round, and as such, the clearing price is predetermined. During the Reveal phase, `RevealClient()` and `RevealMM()` reveal the orders corresponding to `CommitClient()` and `CommitMM()` instances from the previous phase, which are settled in the Resolution phase according to clearing price rules which maximise the amount of notional to be traded, as is in a WSFBA.

When players are rational however, running these algorithms might not be SINCE (Section 3.1). Only if a protocol is SINCE can we be sure that rational players correctly follow the protocol. Towards proving FairTraDEX is SINCE for all rational clients and MMs, we prove that some player in the blockchain protocol runs a `Resolution()` instance every round in Lemma 5.2. Then, in Lemma 5.3 (resp. 5.4) we prove that given a rational client (resp. MM) runs an instance of `Register()` (resp. `CommitMM()`), that same player correctly runs `CommitClient()` and `RevealClient()` (resp. `RevealMM()`) in the proceeding phases. After this, we show in Lemma 5.5 (resp. 5.6) that it is indeed SINCE for a client (resp. MM) to run `Register()` (resp. `CommitMM()`). Composing these results gives us the result we need, as presented in Corollary 5.7, that FairTraDEX is SINCE for clients and MMs.

Lemma 5.2. At least one player runs `Resolution()` in every round.

Lemma 5.3. A rational client who correctly runs an instance of `Register()` also runs correct corresponding instances of `CommitClient()` and `RevealClient()`.

Lemma 5.4. A rational MM who correctly runs an instance of `CommitMM()` also runs a correct instance of `RevealMM()` in the proceeding phase.

Lemmas 5.3 and 5.4 only guarantee that rational clients or MMs who enter the protocol correctly follow the protocol. We also want to ensure rational players join the protocol. i.e. clients run an instance of `Register()`, and MMs run an instance of `CommitMM()`. To show this is the case, we use our results from Section 4.1. We assume a rational player observes no utility (positive or negative) for doing nothing, and has no cost for running the algorithms. In reality, all of these algorithms have costs to run (blockchain transaction fees, proof generation, etc.). FairTraDEX requires these costs to offset against (and be strictly less than) the utility gained by clients (by definition in Section 3) for participation, and MMs from buying below and selling above the MIFP. Thus, the choice of blockchain and its transaction fees dictate the type of clients who can avail of the guarantees provided by FairTraDEX.

Lemma 5.5. Consider an instance of FairTraDEX between A_{tkn} and B_{tkn} . A rational client P_i with $bal_i(\mathfrak{B}) > escrow_{client}$ runs an instance of `Register()`.

Lemma 5.6. Consider an instance of FairTraDEX between A_{tkn} and B_{tkn} , and at least 1 previously called instance of `Register()`. Any rational MM P_i with $bal_i(\mathfrak{B}) \geq escrow_{MM}$ and $bal_i(A_{tkn}), bal_i(B_{tkn}) > 0$ runs an instance of `CommitMM()`.

With these lemmas in hand, we have it that rational clients and rational MMs correctly execute all algorithms as outlined by FairTraDEX. This can be expressed concisely in the following corollary. In this corollary, and the theorem that follows, we assume clients and MMs satisfy the token-balance requirements as described in Lemmas 5.5 and 5.6 respectively.

Corollary 5.7. FairTraDEX is strong incentive compatible in expectation for rational clients and MMs.

²⁰MMs are also allowed to participate as clients if privacy is a concern. `CommitMM()` provides professionals with a functionality to efficiently provide liquidity in a decentralised setting. It is possible to introduce a `RegisterMM()` function analogous to `Register()`, allowing MMs to relay markets in ZK. We believe this has negligible benefit for professionals who already have price and market-size hidden through the commitment scheme.

Now we are in a position to state the main result regarding FairTraDEX. Namely, that for an instance of FairTraDEX between any two tokens A_{tkn} and B_{tkn} , rational players follow the protocol. Moreover, we show that with at least n_ψ Register() calls, the optimal strategy for a client is to submit market orders, while the optimal strategy for a MM with MIFP $y_{A \rightarrow B}$ is to show a market *bid @ offer* with $bid \approx y_{A \rightarrow B} \approx offer$ in the case where there are at least 2 non-cooperative MMs, and of width $w \leq f_{mcf}$ otherwise.

Theorem 5.8. Consider an instance of FairTraDEX between A_{tkn} and B_{tkn} with MIFP $y_{A \rightarrow B}$ and at least n_ψ previously called instances of Register(). For N non-cooperative MMs, the following strategies form strict Nash equilibria:

- $N = 1$: Clients run Register(), followed by CommitClient() producing market orders of width f_{mcf} . The MM runs CommitMM() producing a market of width at most f_{mcf} with reference price equal to $y_{A \rightarrow B}$ in size Q_{not} . Clients and MMs then run RevealClient() and RevealMM() respectively.
- $N \geq 2$: Clients run Register(), followed by CommitClient() producing market orders of width greater than 1. MMs run CommitMM() producing markets of width 1 with reference price equal to $y_{A \rightarrow B}$ in size of at least Q_{not} . Clients and MMs then run RevealClient() and RevealMM() respectively.

Proof. From Corollary 5.7, we know the running of FairTraDEX is SINCE. Furthermore, from Observation 5.1, we have seen that FairTraDEX with at least n_ψ previously called instances of Register() implements a WSFBA. Given FairTraDEX is a WSFBA, the statement follows by applying Theorem 4.4. \square

Remark 5.9. To minimise expected absolute trade imbalances in a DEX auction, existing protocols, including FairTraDEX, require the hiding/mixing of order-information. Consider how FairTraDEX compares to previous DEXs aimed at ensuring client privacy [3, 13, 17]. In these previous protocols, each order commit reveals the same, and in some cases more information per-order than a Register() call in FairTraDEX. EEV protection guarantees in these previous protocols which require n_ψ orders per auction are achieved in FairTraDEX for every order in every auction when n_ψ players are registered to participate in the protocol. This is an n_ψ factor improvement in EEV protection/block-space requirements per auction.

More than this, these previous protocols face liveness issues when players are concerned about EEV. The first players entering one of these previous protocols must choose to do so without any guarantees of protection against EEV attacks based on information leaked from order commitment (trade direction, identity, trading patterns, etc.). In FairTraDEX, a client wanting ψ -EEV protection merely waits until n_ψ players have registered, while retaining the option to withdraw their tokens at any time.

5.4 Smart Contract Implementation of FairTraDEX

A blockchain-based pseudo-code implementation of FairTraDEX, and code description, are provided in Appendix F. We outline here the key differences between the algorithmic description of Section 5.2, and the blockchain-based implementation in Appendix F. As a blockchain-based implementation under the model of Section 5.1 involves a PKI for message sending, all public algorithm outputs must now be signed using the PKI. These messages must now be included in blockchain transactions, with a transaction fee required to ensure the transaction gets added to the blockchain.

For a player to publish a transaction to a blockchain-based smart contract without revealing her identity, she must utilise a relay. Otherwise, the transaction fee is payable from her account, revealing sensitive information such as trading patterns and account balances. Furthermore, this relay must be rewarded on-chain for relaying the transaction. This reward is added by the client when depositing her escrow, and retrievable by the first relay publishing the transaction to the blockchain. Furthermore all checks, such as those for the previous use of serial numbers in CommitClient(), or the recording of the tightest MM width in Resolution(), are explicitly encoded in the blockchain-based implementation.

5.5 Existence of irrational players and coalitions

When analysing the optimal strategies of rational players in WSFBAs, our results are based on all players being rational and that n_ψ instances of Register() are called. If we consider the presence of irrational players in the system, we can apply the following adjustments:

- **Irrational MM:** In Lemma 4.3, it is shown that the optimal strategy for a MM is to show markets centred around the MIFP. Any other (irrational) strategy must therefore result in reduced expectancy for the MM, and higher expectancy for clients. Therefore, given the presence of irrational MMs, submitting market orders maximises client expectancy (with greater expected utility than in the presence of rational MMs, although with increased variance).
- **Irrational client:** Given the optimal strategy for rational clients is to submit market orders, a irrational client may then submit limit orders. This merely reduces the irrational client’s chance of trading vs. other clients. This would not change the strategy of non-colluding rational MMs, but may have some affect on a monopolistic MM’s interpretation of f_{mcf} .

Furthermore, if less than n_ψ instances of Register() are called, clients resort to submitting limit orders. This can be seen in the proof of Theorem 4.4. In the proof, if clients are not sure that a MM will show a market with reference price equal to the MIFP, the case when less than n_ψ instances of Register() are called, the optimal strategy for clients is to submit limit orders, which only stands to reduce the clients’ probability of trading. As the number of non-cooperative players in FairTraDEX decreases towards two, the guarantees of FairTraDEX approach those of an AMM. However, as client price and order size remain hidden until the counterparty chooses her strategy, and before the clearing price is fixed (end of the Commit phase), FairTraDEX maintains advantages over AMMs against client-based EEV attacks, such as price/order-size specific front-running and selective participation.

6 Conclusion

We provide FairTraDEX, a blockchain-based DEX protocol based on FBAs in which we formally prove the strategies of rational participants have strict Nash equilibria in which all trades occur at the MIFP plus or minus bounded upfront costs (specified market widths) which approach 0 in the presence of non-cooperative MMs. This is an attractive alternative to existing mainstream protocols such as AMMs where rational players effectively and systematically prevent such an equilibrium from happening. Compared to previous blockchain-based attempts to implement EEV-proof DEXs, FairTraDEX is the first to practically allow for indistinguishable client-order submissions by decoupling order submission from escrow deposit and order revelation. The FairTraDEX properties outlined in Section 5.3, and summarised in Remark 5.9, provide important improvements on previous protocols regarding EEV protection, setting a new standard for EEV protection in DEXs.

To do this, we show that FairTraDEX implements a width-sensitive FBA, an FBA variant that benefits from the game-theoretic guarantees of an FBA in a distributed blockchain-based setting. We prove our construction is strong incentive compatible in expectation for all rational participants, and combine these results in Theorem 5.8 which outlines the conditions under which clients and MMs trade at the MIFP. In this equilibrium, all fees are known a-priori, with EEV effectively prevented when enough players have registered to participate in the protocol. Implementing, analysing and optimising FairTraDEX for wide-scale deployment makes for exciting future work.

7 Acknowledgements

Thanks to Padraic O’Donoghue for his helpful feedback, discussions and protocol development throughout the production of this paper. Thanks also to IOHK for its support in producing this paper as part of a thoroughly enjoyable internship.

References

- [1] Avi Asayag, Gad Cohen, Ido Grayevsky, Maya Leshkowitz, Ori Rottenstreich, Ronen Tamari, and David Yakira. Helix: A fair blockchain consensus protocol resistant to ordering manipulation. *IEEE Transactions on Network and Service Management*, 18(2):1584–1597, 2021.

- [2] Niko Baric and Birgit Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In *Proceedings of the 16th Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT'97, page 480–494, Berlin, Heidelberg, 1997. Springer-Verlag.
- [3] Carsten Baum, Bernardo David, and Tore Frederiksen. P2dex: Privacy-preserving decentralized cryptocurrency exchange. In Kazue Sako and Nils Ole Tippenhauer, editors, *Applied Cryptography and Network Security*, pages 163–194, Cham, 2021. Springer International Publishing.
- [4] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474, New York, NY, USA, 2014. IEEE Computer Society.
- [5] Josh Benaloh and Michael de Mare. One-way accumulators: A decentralized alternative to digital signatures. In Tor Hellesest, editor, *Advances in Cryptology — EUROCRYPT '93*, pages 274–285, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.
- [6] Daniel Benarroch, Matteo Campanelli, Dario Fiore, Kobi Gurkan, and Dimitris vKolonelos. Zero-knowledge proofs for set membership: Efficient, succinct, modular. In Nikita Borisov and Claudia Diaz, editors, *Financial Cryptography and Data Security*, pages 393–414, Berlin, Heidelberg, 2021. Springer Berlin Heidelberg.
- [7] Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to iops and stateless blockchains. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 561–586, Cham, 2019. Springer International Publishing.
- [8] Eric Budish, Peter Cramton, and John Shim. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response *. *The Quarterly Journal of Economics*, 130(4):1547–1621, 07 2015.
- [9] Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *Proceedings of the 22nd Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '02, page 61–76, Berlin, Heidelberg, 2002. Springer-Verlag.
- [10] Melissa Chase and Anna Lysyanskaya. On signatures of knowledge. In *Proceedings of the 26th Annual International Conference on Advances in Cryptology*, CRYPTO'06, page 78–96, Berlin, Heidelberg, 2006. Springer-Verlag.
- [11] Jing Chen and Silvio Micali. Algorand: A secure and efficient distributed ledger. *Theoretical Computer Science*, 777:155–183, 2019.
- [12] Michele Ciampi, Muhammad Ishaq, Malik Magdon-Ismael, Rafail Ostrovsky, and Vassilis Zikas. Fairmm: A fast and frontrunning-resistant crypto market-maker. Cryptology ePrint Archive, Report 2021/609, 2021. Retrieved: 02/02/2022.
- [13] Theodoros Constantinides and John Cartlidge. Block auction: A general blockchain protocol for privacy-preserving and verifiable periodic double auctions. In *2021 IEEE International Conference on Blockchain (Blockchain)*, pages 513–520, United States, 2021. IEEE Computer Society.
- [14] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *2015 IEEE Symposium on Security and Privacy*, pages 253–270, United States, 2015. IEEE.
- [15] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. <https://arxiv.org/abs/1904.05234>, 2019. Retrieved: 19/01/2022.
- [16] Sanmay Das and Malik Magdon-Ismael. Adapting to a market shock: Optimal sequential market-making. In *Proceedings of the 21st International Conference on Neural Information Processing Systems*, NIPS'08, page 361–368, Red Hook, NY, USA, 2008. Curran Associates Inc.

- [17] Hisham S. Galal and Amr M. Youssef. Publicly verifiable and secrecy preserving periodic auctions. In Matthew Bernhard, Andrea Bracciali, Lewis Gudgeon, Thomas Haines, Ariah Klages-Mundt, Shin’ichiro Matsuo, Daniel Perez, Massimiliano Sala, and Sam Werner, editors, *Financial Cryptography and Data Security. FC 2021 International Workshops*, pages 348–363, Berlin, Heidelberg, 2021. Springer Berlin Heidelberg.
- [18] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology - EUROCRYPT 2015*, pages 281–310, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.
- [19] Wu S. Gjøsteen K., Raikwar M. Pribank: Confidential blockchain scaling using short commit-and-proof nize argument. In *Topics in Cryptology – CT-RSA 2022. CT-RSA 2022*, Cham, 2022. Springer International Publishing.
- [20] S Goldwasser, S Micali, and C Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC ’85, page 291–304, New York, NY, USA, 1985. Association for Computing Machinery.
- [21] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 305–326, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [22] Kobi Gurkan, Koh Wei Jie, and Barry Whitehat. Community proposal: Semaphore: Zero-knowledge signaling on ethereum, 2020. Retrieved: 25/01/2022.
- [23] Aljosha Judmayer, Nicholas Stifter, Philipp Schindler, and Edgar Weippl. Estimating (miner) extractable value is hard, let’s go shopping!, 2021. <https://ia.cr/2021/1231>.
- [24] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020*, pages 451–480, Cham, 2020. Springer International Publishing.
- [25] 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)*, pages 335–353, United States, 2018. IEEE.
- [26] Conor McMenamin, Vanesa Daza, and Matteo Pontecorvi. *Achieving State Machine Replication without Honest Players*, page 1–14. Association for Computing Machinery, New York, NY, USA, 2021.
- [27] Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *Advances in Cryptology — CRYPTO ’87*, pages 369–378, Berlin, Heidelberg, 1988. Springer Berlin Heidelberg.
- [28] Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In *2013 IEEE Symposium on Security and Privacy*, pages 397–411, United States, 2013. IEEE Computer Society.
- [29] Lan Nguyen. Accumulators from bilinear pairings and applications. In *Proceedings of the 2005 International Conference on Topics in Cryptology, CT-RSA’05*, page 275–292, Berlin, Heidelberg, 2005. Springer-Verlag.
- [30] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, 2007.
- [31] Rafael Pass, Lior Seeman, and abhi shelat. Analysis of the blockchain protocol in asynchronous networks. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017*, pages 643–673, Cham, 2017. Springer International Publishing.
- [32] Yupeng Zhang, Jonathan Katz, and Charalampos Papamanthou. An expressive (zero-knowledge) set accumulator. In *2017 IEEE European Symposium on Security and Privacy (EuroS P)*, pages 158–173, United States, 2017. IEEE.

A Extended Related Work

Estimating (Miner) Extractable Value is Hard [23]

This paper attempts to formalise extractable value and generalise it beyond value extractable by miners. We also believe it is necessary to model the decision of all rational players based on *expected extractable value* (EEV) that can be generated by particular orderings of transactions/blocks by any player in the system, and not just the miner. The approach taken is to consider EEV as the maximum of all non-protocol strategies, with protocols considered secure if the EEV of following the protocol is strictly dominated by following the protocol strategy, which is further formalised in [26]. In our paper, we also consider an additional case of EEV not necessarily considered in [23] which is prevalent in commit-reveal protocols such as [12]. In such protocols, honest behaviour usually involves sending a valid second transaction (the reveal transaction in a commit-reveal protocol), but where players can extract value in expectancy by not sending these transactions. However, we believe the definition of EEV in [23] can be extended to include these attacks. As such, we also move away from the legacy use of MEV, and focus instead on the prevention of the more general EEV.

Publicly Verifiable and Secrecy Preserving Periodic Auctions [17]

This protocol also attempts to implement an FBA, and as such has many similarities to FairTraDEX. As in FairTraDEX, the protocol progresses in rounds of Commit, Reveal and Resolution phases. In [17], there is a designated *operator* who is in charge of settling the auction. Players commit to orders in the Commit phase, as well as providing cryptographic information, which is used to prove correct settlement in the Resolution phase. Unlike FairTraDEX, these commit messages are sent by players directly to the blockchain, revealing identity and trade direction. In the Reveal phase, players encrypt their orders using the operator's public key, and send the encryptions to the operator. In the Resolution phase, the operator then chooses a clearing price which intersects the buy and sell liquidity, maximising the notional to be traded. The operator then publishes a list of all matched orders to the blockchain, along with a range proof which is used to verify the correct execution of orders, while not revealing any information about unexecuted orders other than that already revealed in the commit phase.

Block Auction [13]

This protocol attempts to implement an FBA, and is the most similar to FairTraDEX. It is an improvement on [17], with a direct comparison of the two protocols forming the main basis of the justification of [13]. As in [17], the protocol is overseen by an *operator* who is in charge of receiving orders privately from players and correctly executing the auction. As in FairTraDEX, the protocol progresses in rounds of Commit, Reveal and Resolution phases. In the Commit phase, players commit to orders and publish these commitments to the blockchain. Although not revealing the trade direction as is the case in [17], these commit messages are sent by players directly to the blockchain, and as such reveal identity. In the Reveal phase, players encrypt their orders using the operator's public key, and send these encryptions to the operator. The operator then publishes all executed orders, while revealing nothing about unexecuted orders. The validity of execution depends on all players who submitted orders verifying that their order should not have been executed given the list of executions.

P2DEX [3]

The P2DEX protocol is an off-chain MPC protocol run by servers where players can submit orders to exchange tokens from one blockchain to another (although it also appears applicable to one blockchain with many tokens). The orders are encrypted using a threshold secret-sharing scheme with each server receiving a unique share. The protocol has mechanisms to identify double-spending of player funds sent to the servers, and deviation (failure/ misbehaviour) of servers, as the MPC matching protocol is publicly verifiable. As such, all players in the blockchain can verify that a set of orders have been matched correctly, or some of the servers deviated from the protocol. The exchange depends on all servers participating in a secret-sharing

protocol to match orders, with at least one server being honest/not colluding with other servers. As with [13, 17], an emphasis is placed on not revealing unmatched orders.

FairMM [12]

Clients submit orders to a single (monopolistic) MM in an off-chain Σ -protocol. Clients contact the MM with a size, direction (communicated on-chain) and price (communicated off-chain). If the MM accepts, the MM then publishes the trade on the blockchain, otherwise aborts. Client orders are sequentialized so only one order can be executed at a time, preventing the MM from reordering client orders.

FuturesMEX [25]

The FuturesMEX protocol is an MPC version of a DEX with claims of anonymity. In [25], client token balances are kept privately by owners in an off-chain database, so the protocol has limited applicability to a blockchain-based setting. Furthermore, orders are submitted publicly to all participants before being settled. For smaller clients with less connectivity, this is equivalent to showing the order to more-connected counterparties before it is executed. This is a typical source of EEV in existing blockchain protocols (through front-running attacks), and something which is protected against in FairTraDEX.

B Background on Zero-Knowledge Primitives

Proving membership has been traditionally solved using cryptographic accumulators [5], where the prover P computes a value (the accumulator) and, based on it, a set of short membership proofs that the verifier V can easily verify. Three are the approaches to construct set membership proofs: Merkle trees [27], RSA accumulators [2, 7], and pairing-based accumulators [29, 32].

Each approach has its own benefits for public parameters, accumulator or witness size or need of trusted setup. The exact choice depends on the resource constraints of the system. We direct interested readers to [6] for a nice review of the main features of each of the approaches.

When the prover P does not want to reveal the value of x , the membership proof should not leak any information on the value of x . At a high level, the general approach is to guarantee privacy using zero knowledge proofs. Zero knowledge proofs [20] are powerful cryptographic primitives that allow a prover P to prove knowledge of the truth of some statement without revealing the statement contents, to some honest verifier V who needs to be convinced of the truth of the statement provided by the prover. Special mention for its applicability should be made of zero-knowledge proofs that are also non-interactive, that is, proofs that only depend on the prover’s private information about the statement and publicly available information²¹. As such, proofs do not depend on interaction with the verifier. The main features in a NIZK argument are completeness, soundness, and zero-knowledge. Completeness guarantees that if the statement is true, the prover behaving honestly can convince the verifier that the statement is true, while soundness ensures that a dishonest prover cannot convince an honest verifier. Zero-knowledge maintains that the only information learned by the verifier is that the statement is true. However, in practice it is required that the prover knows a witness for the statement, that is, a zero-knowledge proof of knowledge. In this case, soundness is not enough and is required that a prover cannot produce a valid proof unless she knows a witness, even if the prover has seen an arbitrary number of simulated proofs. This is what is known, as simulation intractability. Furthermore, NIZK arguments are interesting for constructing other cryptographic primitives, such as signature of knowledge (SoK) [10].

In literature, there are several constructions that add the privacy layer using zero knowledge proofs for set membership based on RSA Accumulators or Merkle Trees are [4, 9, 28]. In these works the prover proves statements about values that are committed, that is, they follow what is known as a commit-and-prove zero knowledge proof. More recent approaches propose new commit-and-prove for set-membership based on SNARKs [6, 14] or Bulletproofs [19].

²¹This public information can come in many forms, but in [4, 6, 21, 28], it must be generated honestly in a process known as a *trusted setup*. If a prover knows the private information used to generate public proof parameters, the knowledge extraction property cannot exist.

C Proofs

Lemma 4.3

Proof. Let us define the market in terms of y_{ref} and w as described in Section 3, namely, $\frac{y_{ref}}{\sqrt{w}} @ \sqrt{w}y_{ref}$. In the cases of a client buyer and client seller of the swap, we convert client notional orders into the tokens being sold, mark the trades to their post-trade MIFPs using a symmetric impact coefficient of the token swap δ , then reconvert the tokens into notional.

If the client is a buyer of the swap, the client is selling A_{tkn} , with trade size of $\frac{X\mathfrak{B}}{y_A}$ in A_{tkn} . The trade occurs on the token swap offer of $\sqrt{w}y_{ref}$, resulting in the sale of $\frac{X\mathfrak{B}}{y_A} \frac{1}{\sqrt{w}y_{ref}}$ token B_{tkn} s by the MM. Finally, the B_{tkn} s bought by the client have an expected post-trade per-token value of $\sqrt{\delta}y_B$, while the notional acquired by the MM (A_{tkn} s) has an expected value of $\frac{X\mathfrak{B}}{\sqrt{\delta}}$. This is an expected net profit for the MM measured in notional of:

$$\frac{X\mathfrak{B}}{\sqrt{\delta}} - \frac{X\mathfrak{B}}{y_A \sqrt{w}y_{ref}} \sqrt{\delta}y_B = \frac{X\mathfrak{B}}{\sqrt{\delta}} - \frac{\sqrt{\delta}X\mathfrak{B}}{\sqrt{w}} \frac{y_{A \rightarrow B}}{y_{ref}}. \quad (1)$$

If the client is a seller of the swap, the client is selling B_{tkn} , with trade size of $\frac{X\mathfrak{B}}{y_B}$ in B_{tkn} . The trade occurs on the token swap bid of $\frac{y_{ref}}{\sqrt{w}}$, resulting in the sale of $\frac{X\mathfrak{B}}{y_B} \frac{y_{ref}}{\sqrt{w}}$ token A_{tkn} s by the MM. Finally, the A_{tkn} s bought by the client have an expected post-trade per-token value of $\sqrt{\delta}y_A$, while again, the notional acquired by the MM (B_{tkn} s) has an expected value of $\frac{X\mathfrak{B}}{\sqrt{\delta}}$. This is an expected net profit for the MM of:

$$\frac{X\mathfrak{B}}{\sqrt{\delta}} - \frac{\sqrt{\delta}X\mathfrak{B}}{\sqrt{w}} \frac{y_{ref}}{y_{A \rightarrow B}}. \quad (2)$$

We know the buying and selling of $X\mathfrak{B}$ notional at $y_{A \rightarrow B}$ are equally likely by the definition of an MIFP as a perfectly-informed signal. Therefore, the total expected profit is:

$$X\mathfrak{B} \left(\frac{1}{2} \left(\frac{1}{\sqrt{\delta}} - \frac{\sqrt{\delta}}{\sqrt{w}} \frac{y_{A \rightarrow B}}{y_{ref}} \right) + \frac{1}{2} \left(\frac{1}{\sqrt{\delta}} - \frac{\sqrt{\delta}}{\sqrt{w}} \frac{y_{ref}}{y_{A \rightarrow B}} \right) \right). \quad (3)$$

To find the maximum with respect to y_{ref} , we take the first derivative of this formula, and let it equal to 0:

$$\frac{y_{A \rightarrow B}}{y_{ref}^2} - \frac{1}{y_{A \rightarrow B}} = 0. \quad (4)$$

Solving for y_{ref} gives $y_{ref} = y_{A \rightarrow B}$, which is equivalent to the MM strictly maximizing her expected profits by letting $y_{ref} = y_{A \rightarrow B}$. \square

Lemma 5.2

Proof. Consider a Resolution phase where at least one player has not called a `CommitClient()` or `CommitMM()` instance in the preceding `Commit` phase. This player is indifferent to the settlement of orders, and as such the only payoff for that player by running `Resolution()` correctly is the receipt of $resBounty \in \mathbb{R}^+$.

Consider instead the case where all players in the system called at least one instance of `CommitClient()` and/or `CommitMM()` in the preceding `Commit` phase. In this case, all players have an additional payoff for receipt of the tokens currently locked in the protocol. As at least one of the buyers or sellers of the swap must receive a non-zero amount of tokens, the receipt of which having at worst 0 utility. This, in addition to the receipt of $resBounty$ makes the calling of `Resolution()` positive expectancy for at least one player in the system. \square

Lemma 5.3

Proof. By correctly running Register(), a player deposits $escrow_{client}$. The only way to receive $escrow_{client}$ back is to correctly run a RevealClient() instance, which itself can only be run after having run a CommitClient() instance in the previous phase. By construction, $escrow_{client}$ is greater than any incurable losses by running CommitClient() and RevealClient(), with maximal losses occurring where the client's order is settled for tokens that have 0 notional (price goes to 0). As the client's initial deposit had notional value strictly less than $escrow_{client}$, so must the incurred loss. The result follows. \square

Lemma 5.4

Proof. By correctly running CommitMM(), a player deposits $escrow_{MM}$. The only way to receive $escrow_{MM}$ back is to correctly run a RevealMM() instance in the proceeding. By definition, as the MM's bid and offer has notional value greater than or equal to the total notional in the auction, Q_{not} , so must the incurred loss of not revealing a market. Therefore, players running CommitMM() always correctly run RevealMM(). \square

Lemma 5.5

Proof. Consider such a client P_i with minimum client utility f_{mcf} to exchange one token for another. To execute the swap, P_i must first call Register(). Given P_i calls Register(), we know from Lemma 5.3 that P_i also calls CommitClient() and RevealClient(). We know from Lemma 5.2, Resolution() will be run in every round, meaning P_i either trades or the tokens are returned. Given a trade occurs, P_i realises the utility of trading given the restrictions of the order generated by P_i in CommitClient(), which can be chosen to be any value with positive utility (buy below the MIFP/sell above the MIFP). If no trade occurs, P_i 's realised utility (of 0) does not change. Therefore, P_i runs Register(). \square

Lemma 5.6

Proof. Consider such a player P_i . Given Register() was called by some player P_j , P_i knows P_j must call CommitClient() and RevealClient(). Furthermore, P_i knows the total order size is bounded by Q_{imb} (Section 4.1). Given MIFP $y_{A \rightarrow B}$ and the definition of a MM, there is some market $market \leftarrow (bid@offer)$ at which P_i observes positive utility to trade with P_j . Therefore, P_i submits $market$ to the auction. We must now ensure P_i submits the order by calling CommitMM(). By the calculation of order settlement, if P_i submits $market$ through the necessary Register() and CommitClient() calls and P_i submits a market order with finite requested width, no trade happens. As such, P_i runs CommitMM(). \square

D FairTraDEX Algorithms

We now describe in detail the algorithms which together form FairTraDEX, and then describe the main differences between FairTraDEX and a WSFBA.

- Setup(κ, Q_{not}) \rightarrow [$params, y_A, RegIDs, MMCommits, clientCommits, revealedOrders, revealedMkts$]: For a given cryptographic-security parameter κ , output the necessary public cryptographic and ZK parameters in $params$. Set $y_A \in \mathbb{R}^+$ as the indicative price of token A_{tkn} (used to convert restrictions based on escrows into token amounts). Choose $escrow_{client} \in \mathbb{R}^+$ such that the notional of any client order is bounded by $escrow_{client}$, and $escrow_{MM} \leftarrow c \cdot Q_{not}$, for some $c > 1$, with Q_{not} as described in Section 4. Choose a bounty $resBounty \in \mathbb{R}^+$ to reward players for successfully calling Resolution(). Add $escrow_{client}, escrow_{MM}, Q_{not}, resBounty$ to $params$. Set $RegIDs, clientCommits, MMCommits, revealedOrders, revealedMkts \leftarrow []$.
- Register($params, P_i, RegIDs$) \rightarrow [$(S, r), regID, RegIDs$]: If $\{bal_i(\mathbb{B}) \geq escrow_{client}\}$, set $bal_i(\mathbb{B}) \leftarrow bal_i(\mathbb{B}) - escrow_{client}$. Then, randomly generate $S, r \in \{0, 1\}^{O(\kappa)}$, and compute $regID \leftarrow f_{com}(S, r)$. Add $regID$ to $RegIDs$.

- $\text{CommitClient}(params, regID, RegIDs, S, r, clientCommits, y_A) \rightarrow [(order), \pi, S, com, clientCommits]$: If $\{regID \in RegIDs\}$: select a token $tkn \in \{A_{tkn}, B_{tkn}\}$, a trade price $p \in \mathbb{R}^+$, and a trade size amount $size \in \mathbb{R}^+$. If token $tkn = A_{tkn}$, set $y \leftarrow y_A$. Otherwise, set $y \leftarrow y_A \cdot p$ (used to verify client order is less than the escrow given the indicative price for A_{tkn}). If $\{size \leq bal_i(tkn), \frac{escrow_{client}}{y}\}$: select a minimum trade width $w \geq 1$. Set $order \leftarrow (tkn, size, p, w)$ and $com \leftarrow f_{com}(order)$. Finally, generate the signature of knowledge $\pi \leftarrow \text{NIZKSoK}[com]\{(regID, r) : \text{MemVerify}(RegIDs, regID) = 1 \ \& \ regID = f_{com}(S, r)\}$. If no proof corresponding to S has been computed before, add S to $clientCommits$. Otherwise, output \perp .
- $\text{CommitMM}(params, P_i, MMCommits, y_A) \rightarrow [(market), com, MMCommits]$: If $\{bal_i(\mathfrak{B}) \geq escrow_{MM}\}$: select a bid size $size_{bid} \in \mathbb{R}^+$, offer size $size_{offer} \in \mathbb{R}^+$ and prices $0 \leq bid \leq offer$ such that $Q_{not} \leq size_{bid} \cdot y_A \leq escrow_{MM}, (bal_i(A_{tkn}) \cdot y_A) \wedge Q_{not} \leq size_{offer} \cdot y_A \cdot offer \leq escrow_{MM}, (bal_i(B_{tkn}) \cdot y_A \cdot offer)$ (ensures the bid, offer prices, and sizes fall within the bounds of the minimum size required by a WSFBA and the escrow). Set $market \leftarrow (bid, size_{bid}, offer, size_{offer})$ and compute $com \leftarrow f_{com}(market)$. Finally, set $bal_i(\mathfrak{B}) \leftarrow bal_i(\mathfrak{B}) - escrow_{MM}$ and add com to $MMCommits$. Otherwise, output \perp .
- $\text{RevealClient}(params, \pi, S, r, order \leftarrow (tkn, size, p, w, y_A), clientCommits, revealedOrders, P_i) \rightarrow [revealedOrders]$: If token $tkn = A_{tkn}$, set $y \leftarrow y_A$. Otherwise, set $y \leftarrow y_A \cdot p$. If $\{\pi \in clientCommits \wedge \pi = \text{NIZKSoK}[f_{com}(order)]\{(regID, r) : \text{MemVerify}(RegIDs, regID) = 1 \ \& \ regID = f_{com}(S, r)\} \wedge size \leq bal_i(tkn), \frac{escrow_{client}}{y}\}$ (Repeat checks from $\text{CommitClient}()$): set $bal_i(tkn) \leftarrow bal_i(tkn) - size$, $bal_i(\mathfrak{B}) \leftarrow bal_i(\mathfrak{B}) + escrow_{client}$ (return escrow), and add $order$ to $revealedOrders$. Otherwise, output \perp .
- $\text{RevealMM}(params, com, market \leftarrow (bid, size_{bid}, offer, size_{offer}), y_A, MMCommits, revealedMkts, P_i) \rightarrow [revealedOrders]$: If $\{com \in MMCommits, com = f_{com}(market) \wedge Q_{not} \leq size_{bid} \cdot y_A \leq escrow_{MM}, (bal_i(A_{tkn}) \cdot y_A) \wedge Q_{not} \leq size_{offer} \cdot y_A \cdot offer \leq escrow_{MM}, (bal_i(B_{tkn}) \cdot y_A \cdot offer)\}$ (Repeat checks from $\text{CommitMM}()$): set $bal_i(A_{tkn}) \leftarrow bal_i(A_{tkn}) - size_{bid}$, $bal_i(B_{tkn}) \leftarrow bal_i(B_{tkn}) - size_{offer}$, $bal_i(\mathfrak{B}) \leftarrow bal_i(\mathfrak{B}) + escrow_{MM}$ (return escrow), and add the orders $(A_{tkn}, size_{bid}, bid, any), (B_{tkn}, size_{offer}, offer, any)$ to $revealedMkts$. Otherwise, output \perp .
- $\text{Resolution}(params, revealedOrders, revealedMkts, clientCommits, MMCommits, P_i) \rightarrow [CP, imbalance, w_{tight} clientCommits, MMCommits, revealedOrders, revealedMkts]$: If this is not the first time $\text{Resolution}()$ was called, output \perp . Otherwise, calculate the tightest market width w_{tight} of a market in $revealedMkts$. Remove all markets in $revealedMkts$ except one with width equal to w_{tight} , chosen using $revealedMkts$ as a random seed to $f_{com}()$ ²². Remove all orders in $revealedOrders$ with requested width greater than w_{tight} . Calculate the clearing price CP which first maximises notional of orders traded from $revealedOrders \cup revealedMkts$, and then minimises the imbalance $imbalance$.
If CP does not maximise notional traded, and then minimise imbalance, output \perp . Otherwise, if $imbalance > 0$ at CP , pro-rate all buy orders with the lowest bid price above CP based on order size. If $imbalance < 0$, pro-rate all sell orders with the highest offer price below CP based on order size. Then, settle all other buy orders with price greater than or equal to CP , and all other sell orders with price less than or equal to CP . Set $clientCommits, MMCommits, revealedOrders, revealedMkts \leftarrow []$, and $bal_i(\mathfrak{B}) \leftarrow bal_i(\mathfrak{B}) + resBounty$.

D.1 FairTraDEX vs. WSFBA

The main differences between FairTraDEX and a WSFBA are as follows:

- Escrows are used to enforce the correct revelation of players who commit to orders or markets. Escrows are only returned to players if orders are revealed and correspond to a valid commit. Furthermore, escrows are chosen large enough to ensure the reclamation of escrows has strictly higher utility than not, ensuring rational players follow the protocol.

²²Given all markets are revealed, the final value of $revealedMkts$, and as such $f_{com}(revealedMkts|*)$, is unpredictable in the presence of two or more non-cooperative MMs. We prove in Lemma 5.4 that all MMs running $\text{CommitMM}()$ also run $\text{RevealMM}()$. The blockchain based implementation of this function is described in App. F.3

- An algorithm involving deposits and/or withdrawals updates the set of balances for all players, identifying the player calling the algorithm.
- FairTraDEX separates the depositing of client escrow and client order commitments. This is a key functionality necessary to preserve client anonymity and the guarantees of a WSFBA. If a client deposits an escrow in the same instance as committing to an order, that information can be used to identify the player, and imply information about the player’s order. By separating the two, commitment does not require the update of global variables that can be used to identify the client.
- Set-membership proofs in ZK in the CommitClient() algorithm are used to prove that a player committing to a client order has deposited a client escrow. As FairTraDEX separates the deposit and commitment steps, these proofs allow a client who deposited an escrow to generate one (and only one, as ZK proofs reveal S) order per escrow, while only revealing that the order corresponds to a deposited escrow. As the number of deposited escrows increases, the probability that an order commitment matches any particular escrow approaches 0. This replicates the anonymous order submission of a WSFBA.
- Tokenised incentives are used to ensure some player in the blockchain calculates the clearing price, and settles orders correctly.

E Clearing Price Verification

The protocol in Algorithm 4 checks that a given clearing price CP clears the highest notional with respect to A_{tkn} . To do this, it checks the imbalance and total notional that would be settled at CP . Note, that the statements that follow are true given some volume trades at the proposed clearing price, which is asserted in the protocol (line 92).

If the imbalance is 0 (line 97), it can be seen that this price must maximise the volume trade while minimising the absolute value of the imbalance. Higher prices reduces the buying notional/ strictly increases the selling notional, while lower prices strictly reduces the selling notional/increases the buying notional, which creates an imbalance without increasing the notional traded.

If the imbalance is positive (line 99), this implies there will be buy orders (partially) unfilled at or above CP . We can see that the only prices at which more notional can trade must be greater than CP . Thus, the contract checks the next price increment above (line 100), and verifies the total notional traded at that price is less than at CP , or equal, but with a larger absolute imbalance (line 109). If the notional traded is lower at this higher price, the clearing price is correct as a lower price reduces the value of the selling notional, and increases the value of the buying notional. If the notional traded is the same, but the absolute value of the imbalance is higher, this imbalance must be a sell imbalance by the same reasoning (buying notional decreases, selling notional increases). If the absolute value of the imbalance is higher (although negative), the imbalances at all price points above that price are increasing (buying notional decreases, selling notional increases).

The same holds for sell imbalances at a proposed clearing price (line 105). Checking the price point below (line 106) and ensuring the notional is lower, or that the notional traded is the same but with a large absolute imbalance (line 109) guarantees that the proposed clearing price is valid.

F Encoding of FairTraDEX

The following is an overview of the blockchain-based encoding of FairTraDEX, as described algorithmically in Section 5 and encoded in Appendix F.5.

For an arbitrary bit-string $m \in \{0, 1\}^*$, $relay(m)$ indicates broadcasting m to the relay transaction mempool, where m is included as a transaction in the blockchain if and only if it gives the including relayer access to a relayer fee f_r .

F.1 Register

Clients randomly generate S , $r \in \{0, 1\}^{n(\kappa)}$, and compute $regID \leftarrow f_{com}(S, r)$. Then, the client sends a $\langle CLIENT-REGISTER, regID \rangle$ to the blockchain using the blockchain's PKI. Upon addition to the blockchain, this deposits an escrow of $escrow_{client}$ and a relayer fee of f_r to the blockchain (line 22), with $regID$ added to $clients$ (line 23).

F.2 Commit

A client wishing to submit an order of the form $order \leftarrow (tkn, size, p, w)$, first generates a commitment to the order $com \leftarrow f_{com}(order)$. The client then generates a NIZK signature of knowledge $\pi \leftarrow \text{NIZKSoK}(com) \{ (regID, r) : \text{MemVerify}(RegIDs, regID) = 1 \ \& \ regID = f_{com}(S, r) \}$ on the commitment. The client then relays a message of the form $\langle COMMIT, com, S, \pi \rangle$ to the relayer mempool, which is then sent to the smart contract by a relayer. The transaction is only valid if $\text{Verify}(\pi)$ returns 1, and as such, a relayer cannot tamper with com . The contract first checks that the maximum auction notional has not been reached ($currAucNotional < Q_{not}$, line 24).

Furthermore, a valid $\langle COMMIT, com, S, \pi \rangle$ message must not reveal a serial number S which has previously been added to $blacklistedSNs$ (initialised line 18). Serial numbers in $blacklistedSNs$ correspond to client commits that were not revealed according to the protocol during a previous Reveal phase. The escrow corresponding to serial numbers of $blacklistedSNs$ are effectively burned by the protocol, with clients permanently losing access to them. If S is not in $blacklistedSNs$, the order commitment is recorded in $clientCommits$ (line 26), and the relayer who relays the transaction to the blockchain receives the fee (line 27).

MMs who wish to participate generate a market of the form $market \leftarrow (bid, size_{bid}, offer, size_{offer})$, and submit a transaction directly to the blockchain of the form $\langle COMMIT, f_{com}(market) \rangle$. This transaction deposits an escrow of $escrow_{MM}$ to the smart contract. The commitment is then recorded in $MMCommits$ (line 30).

Client and MM Commit transactions are collected until the Commit phase deadline, $requestDeadline$ (line 11), has passed (line 31).

F.3 Reveal

A client who committed to trade through a $\langle COMMIT, com, S, \pi \rangle$ transaction in the Commit phase submits a Reveal transaction directly to the blockchain of the form $\langle CLIENT-REVEAL, tkn, size, p, w, S, r, regID, regIDNew \rangle$ (line 34). If the client intends to re-enter the protocol as a client, $regIDNew$ is a commitment to a new serial number and randomness. Otherwise, it is the null value.

This $\langle CLIENT-REVEAL, * \rangle$ transaction reveals the token being sold, token amount to sell, and requested width of the order. The p either reveals a limit price at which the client is selling, that the order is the market order if $p = mkt$, or that the client is withdrawing their escrow if $p = withdraw$. The contract checks:

- $S \in clientCommits$ to verify a commitment corresponding to that serial number has been recorded.
- $regID = f_{com}(S, r)$ to ensure that client was indeed the same player that generated the $regID$ and that the client order is the same as that committed in the Commit phase $f_{com}(tkn, size, p, w) = clientCommits[S].com$.

If the transaction is valid, the order is added to $revealedBuyOrders$ or $revealedSellOrders$, depending on direction. If the token being sold by the client is A_{tkn} , the effective order size for clearing price calculation and trade size allocation is the minimum of $size$ and $escrow_{client}/y_A$, the maximum token A_{tkn} order size allowable (line 37), with the order recorded in $revealedOrders$ (line 38). If the token being sold by the client is B_{tkn} , the order size is the minimum of $size$ and $escrow_{client}/(y_A \cdot offer)$ (line 45), with the order recorded in $revealedOrders$ (line 46).

Finally, if $regIDNew$ is the null value, the escrow is returned (line 41 or 48), while if it is not, it corresponds to re-entering the protocol as a new client (saving on an additional transaction to re-enter as a client).

A MM who committed to a market through a $\langle COMMIT, com \rangle$ transaction in the commit phase submits $\langle MM-REVEAL, market \leftarrow (bid, size_{bid}, offer, size_{offer}) \rangle$ (line 56). The contract verifies:

- The market $market$ matches the previously commitment from the commit phase ($f_{com}(market) = MMCommits[MM].com$) (line 56).
- For the bid, $Q_{not}/y_A \leq size_{bid}$, which verifies the MM has provided the minimum notional required by a WSFBA (line 57).
- For the offer, $Q_{not}/(y_A \cdot offer) \leq size_{offer}$, which again verifies the MM has provided the minimum notional required by a WSFBA (line 57).

Any MM not revealing a market (line 84) loses their escrow (line 85) and is prevented from participating in the Resolution phase. Otherwise, from the set of all valid revealed markets $revealedMkts$, the tightest market is selected, including a tie-breaking procedure for more than one market with width equal to the tightest width. The tie-breaker used in our implementation of FairTraDEX takes, for a MM MM (identified by a unique public key) and submitted market $market$, the market corresponding to the largest value of $f_{com}(f_{com}(revealedMkts) || MM || market)$ (lines 64-80). Given the tightest market $market \leftarrow (bid, size_{bid}, offer, size_{offer})$ after tie-breaks, the two implicit limit orders are added to the set of revealed orders $revealedBuyOrders$ and $revealedSellOrders$. As was the case with client orders, the effective order size for clearing price calculation and trade size allocation of the bid is the minimum of $size_{bid}$ and $escrow_{MM}/y_A$, while the effective offer size is the minimum of $size_{offer}$ and $escrow_{MM}/(y_A \cdot offer)$.

Reveal transactions are collected until the Reveal deadline, $revealDeadline$ (line 11), has passed (line 31).

F.4 Resolution

Once the protocol enters the Resolution phase, any player in the system can propose a clearing price by submitting a $\langle CP, * \rangle$ message. Players submitting such a message must deposit a token amount (which we set as $resBounty$, although any significantly large value to prevent invalid calls to the smart contract will do). This deposit, along with a bounty is returned to the player if CP is a valid clearing price.

The orders are then settled based on the clearing price CP (lines 114-134). If the quantity of A_{tkn} being sold is greater than the quantity being bought, the sell orders at the highest sell price below the clearing price are pro-rated based on the quantity of A_{tkn} being sold. If the quantity of token A_{tkn} being bought is greater than the quantity being sold, the buy orders at the lowest buy price above the clearing price are pro-rated based on the quantity of A_{tkn} being bought. Remaining unexecuted order balances and escrows are returned to the owners.

F.5 Protocol Encoding

In the following, for arrays containing array objects, the array objects are uniquely identifiable by the first item in the array (i.e. client identifier, serial number, ZKProof).

Algorithm 1 Register

```
1:  $players \leftarrow generatePopulation()$ 
2:  $RegIDs \leftarrow []$ 
3:  $clientCommits \leftarrow []$ 
4:  $MMCommits \leftarrow []$ 
5:  $phase \leftarrow$ 
6:  $w_{tight} \leftarrow any$ 
7:  $revealedBuyOrders, revealedSellOrders, revealedMkts \leftarrow []$ 
8:  $lastPhaseChange \leftarrow 0$  ▷ tracks the block number of last step update
9:  $f_r \leftarrow getRelayerFee()$ 
10:  $minTickSize \leftarrow getMinimumTickSize()$ 
11:  $requestDeadline, revealDeadline \leftarrow T$  ▷ Deadline for responses equal to the maximal reveal delay described in the Threat Model
12:  $escrow_{client} \leftarrow Y$  ▷ Escrow required to show each market, in line with the Threat Model
13:  $Q_{not} \leftarrow getMaxAuctionNotional()$ 
14:  $c \leftarrow random(\mathbb{R}_{>1})$ 
15:  $escrow_{MM} \leftarrow c \cdot Q_{not}$  ▷ Escrow required to show each market, some amount greater than  $Q_{not}$ 
16:  $y_A \leftarrow getTokenAIndicativePrice()$ 
17:  $currAucNotional \leftarrow 0$ 
18:  $blacklistedSNs \leftarrow []$  ▷ Tracks revealed serial numbers that misbehaved

19: function  $Initialise()$ 
20:    $phase \leftarrow Commit$ 

21: upon  $\langle CLIENT-REGISTER, regID \rangle$  from  $player \in players$  with  $player.balance(\mathbb{B}) > escrow_{client} + f_r$  do ▷ register player as a client
22:    $player.transfer(escrow_{client} + f_r, \mathbb{B}, protocolContract)$  ▷ Add client deposit to the contract account
23:    $clients.append(regID)$ 
```

Algorithm 2 Commit

```
24: upon  $relay(\langle COMMIT, com, S, \pi \rangle)$  from  $player \in players$  with  $currAucNotional < Q_{not} \wedge Verify(\pi, com) = 1$   
    $\wedge phase = Commit \wedge \neg(S \in blacklistedSNs)$  do
25:    $currAucNotional \leftarrow currAucNotional + escrow_{client}$ 
26:    $clientCommits.append([S, com])$ 
27:    $protocolContract.transfer(f_r, \mathbb{B}, player)$  ▷ Reward relayer

28: upon  $\langle COMMIT, com \rangle$  from  $player \in players$  with  $player.balance(\mathbb{B}) > escrow_{MM} \wedge \neg(player \in MMCommits)$   $phase = Commit$   
   ▷ Allow only one market per player address
do
29:    $player.transfer(escrow_{MM}, \mathbb{B}, protocolContract)$  ▷ Transfer escrow to the protocol contract account
30:    $MMCommits.append([player, com])$ 

31: upon  $phase = Commit \wedge Blockchain.height() = lastPhaseChange + requestDeadline$  do
32:    $phase \leftarrow Reveal$ 
33:    $lastPhaseChange \leftarrow Blockchain.height()$ 
```

Algorithm 3 Reveal

```

34: upon  $\langle CLIENT-REVEAL, tkn, size, p, w, S, r, regID, regIDNew \rangle$  from  $player \in players$  with  $S \in clientCommits$ 
     $\wedge regID = f_{com}(S, r) \wedge f_{com}(tkn, size, p) = clientCommits[S].com \wedge phase = Reveal$  do
35:   if  $p \neq withdraw$  then
36:     if  $tkn = A_{tkn} \wedge player.balance(A_{tkn}) \geq size$  then
37:        $size \leftarrow \text{minimum}(size, escrow_{client}/y_A)$ 
38:        $revealedBuyOrders.append([player, size, p, w])$  ▷ Add client order to array of orders to trade
39:        $player.transfer(size, A_{tkn}, protocolContract)$ 
40:       if  $regIDNew = \emptyset$  then
41:          $protocolContract.transfer(escrow_{client}, \mathfrak{B}, player)$ 
42:       else
43:          $clients.append(regIDNew)$ 
44:       else if  $tkn = B_{tkn} \wedge player.balance(B_{tkn}) \geq size$  then
45:          $size \leftarrow \text{minimum}(size, escrow_{client}/(y_A \cdot p))$ 
46:          $revealedSellOrders.append([player, size, p, w])$  ▷ Add client order to array of orders to trade
47:         if  $regIDNew = \emptyset$  then
48:            $protocolContract.transfer(escrow_{client}, \mathfrak{B}, player)$ 
49:         else if  $player.balance(\mathfrak{B}) > f_r$  then
50:            $player.transfer(f_r, \mathfrak{B}, protocolContract)$ 
51:            $clients.append(regIDNew)$ 
52:       else ▷ Client wants to withdraw
53:          $protocolContract.transfer(escrow_{client}, \mathfrak{B}, player)$ 
54:          $clients.remove(regID)$ 
55:          $clientCommits.remove(S)$ 

56: upon  $\langle MM-REVEAL, market \leftarrow (bid, size_{bid}, offer, size_{offer}) \rangle$  from  $MM \in MMCommits$  with  $f_{com}(market) = MMCommits[MM].com$ 
     $\wedge phase = Reveal$  do
57:   if  $(Q_{not}/(y_A \cdot offer) \leq size_{offer} \leq player.balance(B_{tkn})) \wedge (Q_{not}/y_A \leq size_{bid} \leq player.balance(A_{tkn}))$  then ▷ Check MM has
    provided the minimum required liquidity,  $Q_{not}$ 
58:      $revealedMkts.append(MM, market)$ 
59:      $MMCommits.remove(MM)$ 

60: upon  $phase = Reveal \wedge len(MMCommits) = 0 \wedge len(clientCommits) = 0$  do ▷ All reveals published
61:    $phase \leftarrow Resolution$ 
62:    $lastPhaseChange \leftarrow Blockchain.height()$ 

63: upon  $phase = Reveal \wedge Blockchain.height() = lastPhaseChange + revealDeadline$  do
64:    $tieBreaker \leftarrow 0$ 
65:    $tightMkt \leftarrow ()$ 
66:    $tieBreakSeed \leftarrow f_{com}(revealedMkts)$  ▷ Generate seed for tie-breaks before revealed markets is changed
67:   for  $MM \in revealedMkts$  do ▷ Select the unique market corresponding to the tie-breaker in the proceeding If statement
68:     if  $(Q_{not}/(y_A \cdot MM.offer) \leq MM.size_{offer} \leq MM.balance(B_{tkn})) \wedge (Q_{not}/y_A \leq MM.size_{bid} \leq MM.balance(A_{tkn}))$  then ▷
    Check MM still has provided the minimum required liquidity
69:      $protocolContract.transfer(escrow_{MM}, \mathfrak{B}, MM)$ 
70:     if  $w_{tight} = any \vee (\frac{offer}{bid} < w_{tight}) \vee (\frac{offer}{bid} = w_{tight} \wedge f_{com}(tieBreakSeed || MM || MM.market) > tieBreaker)$  then
71:        $w_{tight} \leftarrow \frac{offer}{bid}$ 
72:        $tieBreaker \leftarrow f_{com}(tieBreakSeed || MM || MM.market)$ 
73:        $tightMkt \leftarrow [MM, market]$ 
74:        $revealedMkts.remove(MM)$ 
75:      $size_{bid} \leftarrow \text{minimum}(tightMkt.size_{bid}, escrow_{MM}/y_A)$ 
76:      $size_{offer} \leftarrow \text{minimum}(tightMkt.size_{offer}, escrow_{MM}/(y_A \cdot tightMkt.offer))$ 
77:      $revealedBuyOrders.append([player \leftarrow tightMkt.MM, size \leftarrow size_{bid}, p \leftarrow tightMkt.bid, w \leftarrow any])$  ▷ Add tightest market to set
of orders to be settled
78:      $revealedSellOrders.append([player \leftarrow tightMkt.MM, size \leftarrow size_{offer}, p \leftarrow tightMkt.offer, w \leftarrow any])$ 
79:      $tightMkt.MM.transfer(size_{bid}, A_{tkn}, protocolContract)$ 
80:      $tightMkt.MM.transfer(size_{offer}, B_{tkn}, protocolContract)$ 
81:     for  $S \in clientCommits$  do ▷ Add all clients who did not reveal order to blacklist, preventing further commitments
82:        $blacklistedSNs.append(S)$ 
83:        $clientCommits.remove(S)$ 
84:     for  $MM \in MMCommits$  do ▷ MMs who did not reveal market in time
85:        $MMCommits.remove(MM)$  ▷ Remove from protocol without adding to revealedOrders, effectively burning escrow
86:      $phase \leftarrow Resolution$ 
87:      $lastPhaseChange \leftarrow Blockchain.height()$ 

```

Algorithm 4 Resolution: Clearing Price Verification

```
88: upon  $\langle CP, volumeSettled, imbalance \rangle$  from  $player \in players$  with  $player.balance(\mathfrak{B}) > resBounty \wedge phase = Resolution$  do  
89:    $player.transfer(resBounty, \mathfrak{B}, protocolContract)$   $\triangleright$  To prevent Sybil attacks, player must deposit funds which are returned if  
    $CP$  is valid  
90:    $revealedSellOrders \leftarrow revealedSellOrders[revealedSellOrders..width() > w_{tight} \vee revealedSellOrders.width() = any]$   $\triangleright$  Remove  
   sell orders that cannot trade due to requested width  
91:    $revealedBuyOrders \leftarrow revealedBuyOrders[revealedBuyOrders..width() > w_{tight} \vee revealedBuyOrders.width() = any]$   
92:   Assert $(volumeSettled > 0 \vee \text{minimum}(revealedSellOrders.p) < \text{maximum}(revealedBuyOrders.p))$   $\triangleright$  If the indicated clearing  
   price is below the lowest offer/above highest bid, all of the preceding checks pass.  
93:    $buyVolume \leftarrow \text{sum}(revealedBuyOrders[revealedBuyOrders.p \geq CP].size)$   
94:    $sellVolume \leftarrow \text{sum}(revealedSellOrders[revealedSellOrders.p \leq CP].size)$   
95:   Assert $(\text{minimum}(buyVolume/CP, sellVolume) = volumeSettled)$   
96:   Assert $((buyVolume/CP) - sellVolume = imbalance)$   
97:   if  $imbalance = 0$  then  $\triangleright$  We are done  
98:      $SettleOrders(CP, buyVolume, sellVolume)$   
99:   if  $imbalance > 0$  then  $\triangleright$  As the auction is bid at  $CP$ , check if next price increment above clears higher volume OR smaller  
   imbalance  
100:     $priceToCheck \leftarrow CP + minTickSize$   
101:     $buyVolumeNew \leftarrow (buyVolume - \text{sum}(revealedBuyOrders[CP \leq revealedBuyOrders.p < priceToCheck].size))/CP$   
102:     $sellVolumeNew \leftarrow sellVolume + \text{sum}(revealedSellOrders[CP < revealedSellOrders.p \leq priceToCheck].size)$   
103:    Assert $((\text{minimum}(buyVolumeNew, sellVolumeNew) < volumeSettled) \vee$   
     $(\text{minimum}(buyVolumeNew, sellVolumeNew) = volumeSettled \wedge imbalance \leq |buyVolumeNew - sellVolumeNew|))$   $\triangleright$  If the next  
    price clears less volume, or clears the same volume with a larger imbalance, the proposed  $CP$  is valid  
104:     $SettleOrders(CP, buyVolume, sellVolume)$   
105:   if  $imbalance < 0$  then  $\triangleright$  As the auction is offered at  $CP$ , check if next price increment below clears higher volume OR  
    smaller imbalance  
106:     $priceToCheck \leftarrow CP - minTickSize$   
107:     $buyVolumeNew \leftarrow (buyVolume + \text{sum}(revealedBuyOrders[CP > revealedBuyOrders.p \geq priceToCheck].size))/CP$   
108:     $sellVolumeNew \leftarrow sellVolume - \text{sum}(revealedSellOrders[CP \geq revealedSellOrders.p > priceToCheck].size)$   
109:    Assert $((\text{minimum}(buyVolumeNew, sellVolumeNew) < volumeSettled) \vee$   
     $(\text{minimum}(buyVolumeNew, sellVolumeNew) = volumeSettled \wedge imbalance \leq |buyVolumeNew - sellVolumeNew|))$   
110:     $SettleOrders(CP, buyVolume, sellVolume)$   
111:    $protocolContract.transfer(2resBounty, \mathfrak{B}, player)$   $\triangleright$  Return deposit, and reward player for submitting valid clearing price
```

Algorithm 5 Resolution: Settle Orders

```
112: function  $SettleOrders(CP, buyVolume, sellVolume)$   
113:    $buyVolume \leftarrow buyVolume/CP$   $\triangleright$  Convert sell volume to equivalent in  $A_{tkn}$   
114:   if  $buyVolume > sellVolume$  then  $\triangleright$  pro-rate buy orders at the min price above (or equal to) the clearing price  
115:      $p_{pro-rate} \leftarrow \text{minimum}(revealedBuyOrders[revealedBuyOrders.p \geq CP].p)$   
116:      $size_{pro-rate} \leftarrow \text{sum}(revealedBuyOrders[revealedBuyOrders.p = p_{pro-rate}].size)/CP$   
117:     for  $order \in revealedBuyOrders[revealedBuyOrders.p = p_{pro-rate}]$  do  
118:        $protocolContract.transfer(order.size \cdot (1 - \frac{buyVolume - sellVolume}{size_{pro-rate}}), A_{tkn}, order.player)$   $\triangleright$  return tokens not going to be  
       exchanged  
119:        $order.size \leftarrow order.size \cdot \frac{buyVolume - sellVolume}{size_{pro-rate}}$   
120:   else if  $sellVolume > buyVolume$  then  $\triangleright$  pro-rate sell orders at the max price below (or equal to) the clearing price  
121:      $p_{pro-rate} \leftarrow \text{maximum}(revealedSellOrders[revealedSellOrders.p \leq CP].p)$   
122:      $size_{pro-rate} \leftarrow \text{sum}(revealedSellOrders[revealedSellOrders.p = p_{pro-rate}].size)$   
123:     for  $order \in revealedSellOrders[revealedSellOrders.p = p_{pro-rate}]$  do  
124:        $protocolContract.transfer(order.size \cdot (1 - \frac{sellVolume - buyVolume}{size_{pro-rate}}), B_{tkn}, order.player)$   $\triangleright$  return tokens not going to be  
       exchanged  
125:        $order.size \leftarrow order.size \cdot \frac{sellVolume - buyVolume}{size_{pro-rate}}$   
126:   for  $order \in revealedBuyOrders, revealedSellOrders$  do  $\triangleright$  iterate through orders  
127:     if  $order \in revealedBuyOrders \wedge (order.p \geq CP \vee order.p = mkt)$  then  $\triangleright$  execute buy order if bid greater than clearing  
     price  
128:        $tokenTradeSize \leftarrow order.size/CP$   
129:        $protocolContract.transfer(tokenTradeSize, B_{tkn}, order.player)$   
130:     else if  $order \in revealedSellOrders \wedge (order.p \leq CP \vee order.p = mkt)$  then  $\triangleright$  execute sell order if bid greater than  
     clearing price  
131:        $tokenTradeSize \leftarrow (order.size)/CP$   
132:        $protocolContract.transfer(tokenTradeSize, A_{tkn}, order.player)$   
133:     else  $\triangleright$  Order not executed  
134:        $protocolContract.transfer(order.size, order.tkn, order.player)$   
135:    $phase \leftarrow Commit$   
136:    $currAucNotional \leftarrow 0$   
137:    $revealedBuyOrders, revealedSellOrders \leftarrow []$   
138:    $w_{tight} \leftarrow any$   
139:    $lastPhaseChange \leftarrow Blockchain.height()$ 
```

G Practical Considerations for FairTraDEX

Escrow choices

Choosing escrow amounts for clients and MMs should reflect the emergent use cases of the protocol. It is possible to create different FairTraDEX instances for the same pair of tokens based on trade size, both for liquidity purposes (MMs will require wider markets for larger-sized orders, but the corresponding increased escrow requirements might prevent smaller clients from participating) and auction use-cases (day-trading vs. end-of-day balancing). Furthermore, as the escrow denomination (\mathfrak{B}) in our description is different to at least one of the tokens, there needs to be some way to translate the escrow amount into order sizes. This will depend on the environment, but on Ethereum for example, price oracles (existing AMMs, Chainlink²³, etc.) can be used. It is also possible to use previous clearing prices from within the FairTraDEX ecosystem, although a self-referential oracle must be implemented carefully as there may be game-theoretic implications. If a satisfactory price oracle exists, deposits can be made in the respective tokens of the swap, further reducing the capital requirements for players and encouraging adoption.

Incentive compatibility given transaction fees

In Section 5, we mention that our Nash equilibria are dependent on the utility gained by clients and MMs being greater than the cost for participation. The choice of smart-contract enabled blockchain on which to deploy will dictate the barrier of entry for clients and MMs alike. Like existing attempts to implement blockchain-based FBAs [13, 17], we have an amortised number of transactions per player of two. A naive comparison to AMMs, where this is reduced to 1 for clients, and 0 for MMs, certainly has less direct costs than FairTraDEX. However, when factors like impermanent loss, slippage²⁴, front-running, and EEV attacks in general, the value being extracted from DEXs incurs a significant indirect cost for clients. Immediately, we can increase the expected cost of using AMMs by the slippage required by AMMs (set to 0.5% as of writing in Uniswap V3, but for larger orders this must increase by the nature of AMMs). We can increase this further by the probability orders are not executed (where prices move more than the slippage, potentially due to front-running) but are added to the blockchain. As such, the indirect costs are substantial, are increasing in order-size increases and proportionally to improvements in client trading ability as strategies can be replicated/front-run. A thorough comparison of the monetary costs of FairTraDEX vs. AMMs over various order-sizes, and trading scenarios makes for interesting future work as FairTraDEX begins to be deployed and tested in the wild.

It can be seen from the proof of Theorem 4.4 that the client strategies identified are strong incentive compatible in expectation as the MM always shows markets of width less than or equal to f_{mcf} . However, the MM strategy of showing width 1 markets is not strong incentive compatible. In addition to the fees described in the protocol, an additional fee can be applied within the protocol itself to incentivise the participation of MMs. This can be a function of MM participation/market widths. As with all additional fees/rewards, the game-theoretic implications of such an incentivisation scheme must be considered.

In our encoding of FairTraDEX, we do not explicitly introduce a cost for clients and MMs in Commit/Reveal phases to reward the submission of the clearing price. In reality, the result in Lemma 5.2 holds without the introduction of an explicit reward, as there all participating clients and MMs will have positive expectancy to receive tokens through correct order resolution. The use of an explicit reward is for illustrative purposes, and to avoid complications regarding transaction fees for running the clearing price checks. The costs of running the Resolution contracts must be ensured to be less than the utility gained by at least one player in the blockchain protocol for calling the contract.

²³<https://chain.link/>

²⁴<https://docs.uniswap.org/protocol/concepts/V3-overview/swaps>