KRYPTON

MICHAEL CHRISTOPH NOWOTNY

Abstract. Krypton is a novel peer-to-peer DeFi exchange protocol that solves the issue of miner extractable value in the context of trading. It is fully decentralized, resistant to front-running and to short-term adverse selection. This is achieved by executing orders as continuous streams over time at finite trading speeds in contrast to concentrating trade execution to a distinct instant of time at infinite speed.

The protocol features price discovery. Rather than relating price to quantity, traders submit demand and supply schedules relating price to trading speed. This change in perspective affords enormous conceptual clarity, predictability, and simplicity: The market clearing equilibrium is simply defined as the intersection of the aggregate demand and supply curves in price/trading-speed space.

An actor in possession of short-term information that needs to be monetized prior to becoming public will need to request a high trading speed, which is limited by the maximum aggregate trading speed offered by the opposite side and increases the price. The unilateral price and trading speed jump induced by a fast trade signals the presence of asymmetric information and gives uninformed traders time to respond by adjusting their quotes. In infinite speed trading systems such as automated market makers (AMMs) and central limit order books (CLOBs), the maximum tradable quantity is exchanged in the very same moment the order that signals the presence of superior information arrives, making it too late for everyone else to react.

Date: December 15, 2021.

Contents

1. Introduction

Front-running and adverse selection of regular traders by faster and more sophisticated actors with short-term information advantages are pervasive in both crypto markets and traditional finance. In centralized markets, the profits arising from these activities are captured by algorithmic high-frequency trading (HFT) firms who compete among each other, devising ever more intricate techniques to uncover opportunities by predicting short-term order flow and to reduce latency of market access to win the race to the centralized exchange against one another. The HFTs' gains from front-running and adverse selection represent costs to regular traders such as pension funds, mutual funds, and retail investors. These costs are the primary reason for the existence of price impact and bid-ask spreads.[1]

Exploiting short-term market microstructure effects does not contribute to efficient price discovery since no new information about a security's fundamentals is uncovered. The cost thus imposed onto regular traders reduces their incentives to trade and hence prevents them from implementing optimal investment strategies. Since these trading activities represent not merely a transfer of wealth from regular traders to sophisticated actors but result in a deadweight loss due to this inefficiency, the activities are collectively known as toxic trading.

Competition among HFTs requires significant investment in research, hard- and software engineering, as well as the construction of networks of radio towers and fiber optic cables for the sole purpose of reducing latency of market access. The fruits of these research and development efforts are kept as trade secrets which leads to duplication of work by every single algorithmic fund in the space. In aggregate, this investment leads to a reduction in economic welfare.

On centralized exchanges such as NYSE, NASDAQ or CME, high-frequency traders are responsible for the majority of the trading volume, producing a significant fraction of total revenue for companies operating these exchanges. Exchange operators have no incentive to implement measures that limit HFT activity as doing so would lead to a substantial reduction in their own revenues. Exchange operators and sophisticated algorithmic traders essentially share the proceeds from toxic trading. This alignment of incentives makes it less likely for centralized incumbents to compete with novel exchange mechanisms in the short term.

In decentralized, trustless and permissionless public blockchain ecosystems, miners[2] are the final arbiters of transaction sequencing, giving them the ability to front-run each trade on every block they win. The privilege to front-run a trade after it has been broadcast to the blockchain network translates into negative latency, turning miners into the ultimate front-runners. The profits that miners can thus appropriate are collectively known as miner extractable value or maximal extractable value (MEV).

[1] Glosten and Milgrom (1985) explain the existence of a bid-ask spreads on information theoretic grounds. Easley and O'Hara (1987) demonstrate that large trade sizes are more likely to be motivated by asymmetric information than small trades and therefore produce higher market impact.

[2] More generally block producers (miners in PoW and validators in PoS).

The explosive growth of decentralized finance (DeFi) from virtually zero in early 2020 to an average daily trading volume of $3M in August 2021 has prompted an equally impressive rise in realized MEV.[3] This rapid growth in MEV has been fueled by Flashbots, a project that distributes a modified Ethereum client called 'mev-geth' that implements an auction mechanism through which algorithmic traders can bid for the privilege to determine the composition of the next block won by a miner running the software.[4] The availability of Flashbots absolves mining operators from developing algorithmic trading sophistication themselves - a factor that would have counteracted the adoption of these practices.[5] As of April 2021, about 84% of miners active on the Ethereum mainnet are running Flashbots' modified mev-geth client, up from 58% in March, 12% in February, and just 4% in January.[6] The rapid rise in Flashbots adoption will ensure that soon almost every Ethereum block produced will be squeezed for its maximal extractable value. The virtual elimination of gas fees on rollups, side chains, and high-performance L1 solutions will make it profitable to front-run even the smallest orders.

The only way to solve this rapidly rising MEV crisis in a way that is compatible with the ideals of Ethereum is to design mechanisms that make front-running worthless or impossible. This is the route we have chosen for Krypton.

In its pursuit of MEV resistance, Krypton switches out the foundational mechanism that existing trading protocols, both in DeFi and in CeFi are based on for a novel approach that solves the pertinent problems of front-running, adverse selection based on short-term information advantages, thus preventing miner extractable value. In Krypton, trades are executed over periods of time rather than being carried out in bursts concentrated to singular points in time which translates to an infinite trading speed. Krypton gives regular traders a way to credibly signal that they are not trying to front-run or exploit short-term information via a choice of trading speed. Regular traders are in no rush to monetize short-lived information and can therefore select a low trading speed. This fundamental change in perspective, which is developed in detail throughout this paper, renders technical superiority and transitory information advantages worthless. By Krypton's design, even the fastest, most sophisticated, and well-informed traders will not be able to extract meaningful value or cause noticeable damage to regular traders - no matter their effort or expense.

The decentralized spot trading protocol for ERC-20 tokens will be extended to derivative contracts in a future version of the protocol that adheres to the same basic principles. Generally speaking, the conceptual approach of addressing the issue of transitory information advantages via continuous trading at limited speeds could be used in applications

[3]Schmidt 2021 estimates MEV to be between $1m and $4m per day.

[4]Flashbots reports total MEV from January through August 2021 to be $549.1M which translates to a daily average of $2.26M during that period.

[5]Proponents of Flashbots argue that the enormous profit potential from MEV would likely have led to an increase in centralization of mining operations since merging a mining pool not engaging in MEV with one that does keeps the efforts to uncover opportunities constant while increasing the probability of winning a block, thus successfully realizing MEV profit.

[6]These numbers have been obtained from Flashbots' monthly and bi-monthly Flashbots transparency reports.

outside of trading such as insurance, sports betting, or other fields where adverse selection matters.

This paper is organized as follows. Section 2 details how toxic order flow is able to exploit regular trades in existing exchange designs. 3 presents the technical approach to implement Krypton using smart contracts. Section 4 concludes.

2. How Toxic Traders Exploit Existing Exchange Mechanisms

2.1. **Adverse Selection in a Central Limit Order Book.** In a central limit order book (CLOB), market makers and large institutional traders place limit orders that sit unexecuted until they are taken by an active trader.

Suppose that a regular trader wants to sell 10,000 shares for at least $100 per share. If an informed trader gains a short-term information advantage by learning that the fair value of the asset has increased to $105, he can swoop in and lift the entire 10,000 shares for $100 before the seller learns of the information and can react.

This has led to a technological arms race in which ever more sophisticated algorithms are devised to either conceal trade intentions via splitting large orders for institutions over time to make them look as random as possible or to reveal them via sophisticated statistical filters in order to front-run.

2.2. **HFTs and Front-Running in a Central Limit Order Book.** Robinhood to predict and exploit short term price movements. Algorithmic high-frequency traders (HFTs) purchase order flow from retail brokers such as

Suppose that an HFT learns that Alice has submitted a market order to purchase 10,000 shares. The current price is $100. The HFT uses its low latency market access to place a buy order and purchase 10,000 shares, then immediately places a sell order for 10,000 shares below the best ask, before Alice's order arrives, selling to Alice at a higher price than they bought at for an immediate profit. If Alice had instead placed a limit order at $101, the HFT would reduce the quantity it transacts to make sure their sell order won't exceed Alice's limit price.

2.3. **Adverse Selection in Automated Market Makers (AMMs).** The AMM creates liquidity by representing a passive side who is willing to act as the counterparty to any trade according to a computationally inexpensive formula. This simple formula loses to arbitrageurs and other toxic traders with short-term information advantages and needs to be compensated for this systematic loss (euphemistically designated 'impermanent loss') via 25 basis points (bps) in fees.

If the the fair price of token X as discovered on two-sided exchanges (such as a CLOB) rises from $100 to $110, an arbitrageur can buy from the AMM until its price has been driven up to $110. This arbitrage mechanism is integral to the AMM as it would otherwise not know the fair price of a token, which would make it an unattractive trading venue. Regular traders are ultimately paying for liquidity provider (LP) losses to informed traders via high trading fees. This is the central inefficiency in peer-to-pool DeFi protocols.

2.4. **Front-Running by Miners in Automated Market Makers.** Decentralized blockchain ecosystems give miners or validators the privilege to choose the ordering of transactions when they produce a block. In their Uniswap V2 audit report, Currin, Terry, Erfurt, Livnev, Manacorda, and Lundfall (2020) highlight the following attack vectors by which miners can extract value from the protocol:

- Order sequencing: A miner wants to sell Token X at an advantageous price and orders all buy transaction to appear before their own sell transaction to drive up the price prior to the sale and places other sell transactions after theirs.
- Sandwich attack: A miner observes that Alice wants to buy a large quantity of Token X, inserts their own purchase driving up the price before recording Alice's original trade, which drives up the price even further, then immediately puts in a sell order at this higher price to profit from this a pump-and-dump scheme.
- Liquidity sandwich attack: A miner observes large trade from Alice which they front-run with massive liquidity provision, giving them majority stake in the LP pool when Alice's order executes, thus securing the lion's share of LP rewards. In the same block, the miner withdraws liquidity immediately after the trade for a riskless profit. This can be implemented using a flash loan.

Until the beginning of 2021, technically sophisticated bot networks and miners used to front-run DeFi trades, liquidations, and withdrawal of AMM liquidity due to extreme price moves. They furthermore engaged in arbitrage and monetized short-term information advantages. The comprehensive adoption of Flashbots' modified mining software in 2021 has shifted these activities onto miners themselves, who are able to exploit privilege to determine transaction order in blocks they produce. Figure 1 illustrates that Flashbots platform, by their own account, has facilitated toxic trading activities on the order of $547M from January through August 2021. This number represents a lower bound since it only includes activities known to Flashbots on selected subset protocols in the DeFi space.

In analogy to the effect of high-frequency trading (HFT) in traditional markets, this has led to a socially wasteful technological arms race among algorithmic traders that degrades UX and impedes the realization of DeFi's promise and potential.

3. The Protocol

3.1. **Technical Challenges.** In Krypton, equilibrium prices can change whenever an order has been completely executed, an order that is currently being executed is modified or canceled, or a new order is placed that trades in equilibrium. Computing the equilibrium entails running the matching engine algorithm, which is so computationally intensive by smart contract standards that it exceeds the block gas limit on Ethereum with more than about 200 orders on each side. Furthermore, transferring tokens every time the equilibrium price changes is inefficient. To reduce the number of on-chain transactions, Krypton collects aggregate token transfers over periods of time during which a number of different market clearing prices may have prevailed and accurately executes those in regular intervals. The technical approach is presented in the remainder of this section.

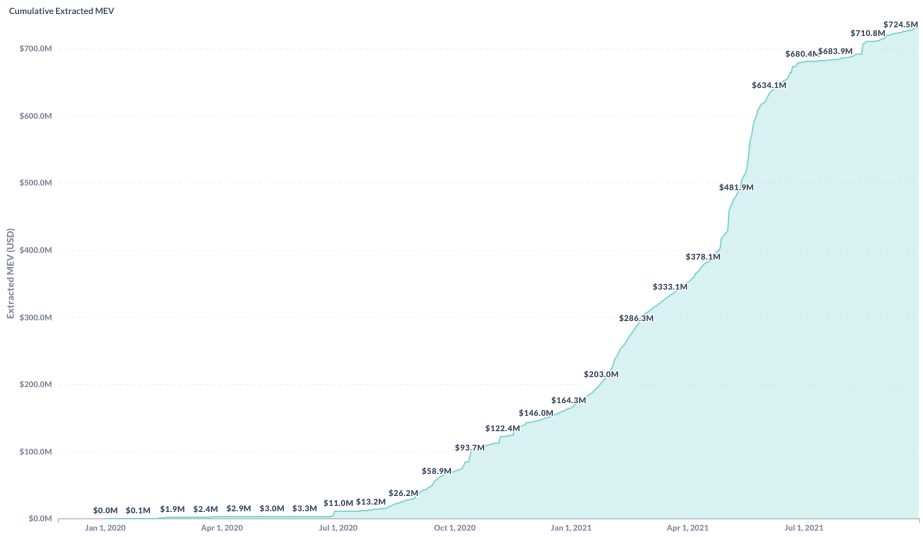

FIGURE 1. Cumulative Historical Miner Extracted Value (MEV). Source: https://explore.Flashbots.net, retrieved on September 29, 2021

3.2. **A Decentralized, Trustless, On-Chain/Off-Chain Approach.** Krypton moves the computationally most intensive parts to a decentralized network of Chainlink oracles via an external adapter. The storage footprint is minimized by only keeping structures of existing and completed orders as well as requests to make, change, or cancel orders on-chain. The computational demand is minimized by solely executing trade settlement on-chain.

Trade execution is run in regular intervals covering the preceding period, consisting of blocks [n, n + h]. "h" could be a minute, an hour, or an entire day for instance. During trade execution, the matching engine is run sequentially and the number of tokens to be transferred for periods between price changes is computed and summed up. Only aggregate transfers are reported back to the smart contract and implemented on-chain. The front-end keeps users up to date regarding the state of the system, anticipating pending changes yet to be memorialized on-chain.

Figure 2 illustrates the relationships between the protocol's participants.

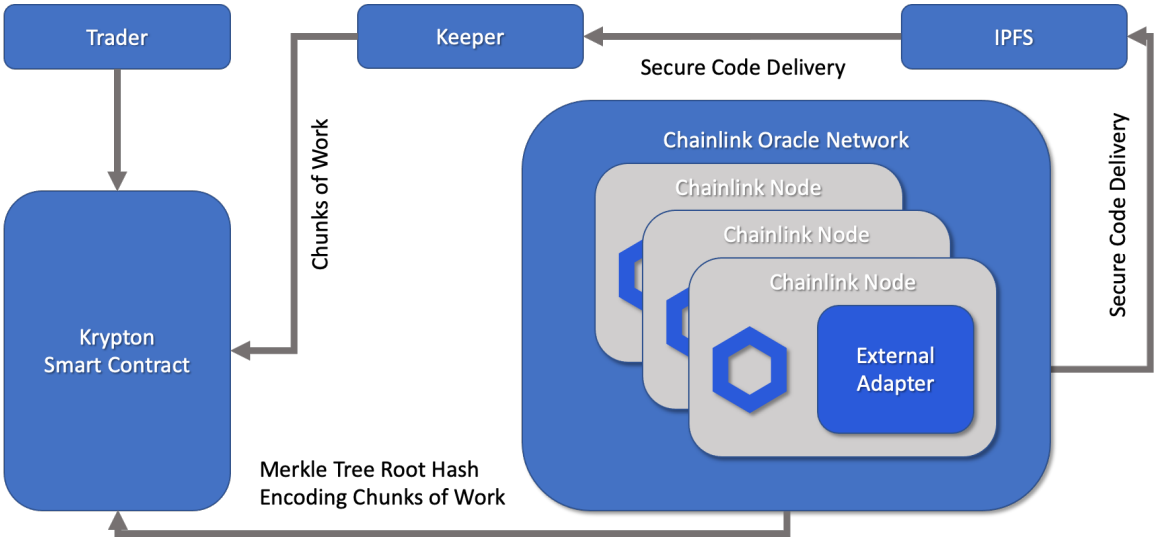

FIGURE 2. Schematic Illustration of Krypton's Protocol Architecture

4. CONCLUSION

Krypton solves information-theoretic inefficiencies in AMMs, central limit order books, and centralized crypto exchanges, as well as traditional exchanges like NYSE or CME. These trading venues are plagued by front-running and monetization of short-term information advantages. These toxic activities prevent the implementation of efficient dynamic trading strategies which result in suboptimal exposure to systematic risk factors as well as unnecessary exposure to idiosyncratic risk.

In decentralized blockchain ecosystems, miners are the final arbiter of transaction sequencing. Their collusion facilitated through Flashbots allows algorithmic traders to implement toxic trading strategies to exploit microstructural characteristics of existing market and protocol mechanisms for the maximum extractable value (MEV) free of unpredictability or risk of interference.

Krypton's approach renders algorithmic trading tools blunt and miners' block sequencing rights worthless. Toxic traders have no way to win on Krypton and will go elsewhere. This will make trading inexpensive and enable of trading strategies that implement an optimal (dynamic) risk-return tradeoff.

References

Currin, D., D. Terry, D. Erfurt, L. Livnev, L. Manacorda, and M. Lundfall (2020): "Uniswap V2 Audit Reeport: Front-Running and Transaction Reordering," .

Easley, D., and M. O'Hara (1987): "Price, trade size, and information in securities markets," *Journal of Financial Economics*, 19(1), 69–90.

Glosten, L. R., and P. R. Milgrom (1985): "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, 14(1), 71–100.