Dex Order Routing

Orchard

Decentralized exchanges (DEXs) are foundational to decentralized finance (DeFi), enabling on-chain trading directly between users without relying on traditional intermediaries. Within the DEX landscape, two primary models dominate: Automated Market Makers (AMMs) and Central Limit Order Books (CLOBs). Automated market makers occupy a central place in decentralised finance as the primary means for on-chain trading. An AMM is a two-sided platform where traders (aka takers) consume the liquidity provided by liquidity providers (LPs) and which operates entirely on-chain.

Liquidity fragmentation remains a major obstacle to efficient cryptocurrency trading. Assets are spread across centralized exchanges (CEXs), decentralized exchanges, and separate blockchain networks. Although many of the existing efforts attempt to improve local efficiency, they do not offer a complete solution to the liquidity fragmentation problem.

1 Introduction

The rise of cryptocurrency trading has transformed global finance by allowing digital assets to move across borders without relying on centralized authorities. However, one of the biggest hurdles the industry faces is the fragmentation of liquidity. In contrast to traditional financial systems where most trading activity is concentrated on a few major platforms, the crypto landscape is highly dispersed. Trading occurs across centralized exchanges, decentralized exchanges, and separate blockchain networks, each with its own isolated pool of liquidity. This decentralized structure makes it challenging for traders to access optimal pricing and sufficient market depth, especially for large transactions. Additionally, the growth of decentralized finance has added layers of complexity, introducing mechanisms like automated market makers that operate independently and often lack synchronization with broader markets.

Institutional investors are particularly affected by fragmented liquidity. For example, when a hedge fund attempts to purchase a significant amount of Ethereum, a single exchange might not have enough supply to fill the entire order without causing a major price spike. To manage this, the fund would likely need to divide the trade across several platforms, which leads to higher costs and increased exposure to market fluctuations. During volatile times, fragmented liquidity can make price swings even more extreme, increasing risks for everyone in the market.

In general, liquidity fragmentation poses a major challenge to the efficiency of cryptocurrency markets. It drives up trading expenses, complicates large-scale transactions, and intensifies volatility-related risks. Addressing this issue is essential for creating a more resilient and reliable crypto trading environment.

Related work

Improving execution performance is, of course, a concern in traditional finance. Guilbaud and Pham [1] emphasize the necessity of order splitting to mitigate the impact of large orders. Later this work was extended to multi-venue routing, focusing on optimising execution cost across varied liquidity profiles. The structure of AMMs inherently adjusts asset prices based on trade sizes. In that respect it is analogous to some traditional models of price impact, such as those incorporating marginal supply-demand curves. Because of this response of price to order size, optimised executions need to split orders across multiple AMMs.

Trade routing plays a crucial role in decentralized exchanges (DEXs). Various DEX platforms employ different algorithms to identify the best trading routes for users. For instance, platforms like Uniswap, PancakeSwap, and SushiSwap utilize algorithms based on depth-first search. Meanwhile, the DEX aggregator 1inch initially implemented a similar approach but later transitioned to a more advanced algorithm, the specifics of which remain undisclosed. Although DFS-based methods are computationally efficient, they don't always guarantee the most profitable trading routes for users. Recently, Zhang, Li, and Tessone [2] improved upon the depth-first search (DFS) graph heuristics most commonly used by modern DEX aggregators by introducing a novel "line-graph-based algorithm". Although it enables more profitable routing, its computational complexity remains too high for practical deployment.

Angeris and Chitra [3] analysed a specific class of constant function AMMs from the perspective of convex optimization theory. Furthermore, Danos, Khalloufi, and Prat [4] framed optimal routing and arbitrage as global convex optimisation problems on arbitrary networks of AMMs. Later, Angeris et al. [5] formulated optimal routing for trades across multiple AMMs, incorporating multiple tokens as inputs within a single optimisation framework. More recently, Diamandis et al. [6] presented an efficient algorithm for optimal routing through constant function market makers using a dual decomposition approach. In contrast, Loesch and Richardson [7] introduced a new framework for optimal routing and arbitrage that moves away from the traditional convex optimization paradigm. Instead of solving a high-dimensional optimization problem, it reframes the problem as a much lower-dimensional root-finding problem.

In summary, all previous algorithms either fail to find optimal solutions or suffer from prohibitive computational complexity. The most practical ones are limited to finding only linear routing paths. While convex optimization has shown the most promise for DEX routing, a naive application of existing algorithms cannot effectively account for network fees and lacks scalability in realistic, practical settings.

Our Proprietary Routing Algorithm

To address these limitations, we have developed a new, proprietary routing algorithm that offers significant advantages over existing solutions. The main advantages of our algorithm are fivefold:

- **AI-Powered Search Space Pruning:** We use our own in-house developed AI model to prune the search space significantly. This allows us to quickly identify the most promising routing options, dramatically reducing the computational overhead and enabling us to find the best routes faster.

- **Scalable Convex Optimization:** We have built upon existing research to ensure that a similar-in-spirit approach to convex optimization can be applied to real-world scenarios with hundreds of thousands of DEXes. Our algorithm is designed to be highly scalable, allowing us to efficiently find the best routes even in the most complex and fragmented DeFi landscapes.

- **Comprehensive Network Fee Handling:** By combining these two advances, our algorithm is able to handle network fees effectively. This ensures that the routes we provide are not only optimal in theory but also in practice, delivering the best possible execution price for our users.

- **GPU Acceleration:** Our algorithm is designed to be easily scaled on GPUs. This architecture allows us to take full advantage of the massive parallel processing capabilities and recent performance improvements in GPU hardware, ensuring our routing calculations are performed at unparalleled speeds. GPUs are leveraged in two distinct ways within our solution: for the training of our AI-powered search space pruning model, and for the actual execution (inference) of the routing algorithm itself. The training phase for our AI is not as computationally expensive as in mainstream large language model (LLM) training and can be completed efficiently on a single GPU. Furthermore, the inference phase for the routing algorithm is highly optimized, allowing a single modern GPU to serve numerous users simultaneously, significantly enhancing overall efficiency and scalability.

- **Certificate of Optimality:** In certain cases, when the AI-driven pruning heuristic preserves favorable geometric properties of the solution space, our algorithm can provide a formal certificate of optimality. This certificate mathematically proves that the found solution satisfies the Karush-Kuhn-Tucker (KKT) conditions for the complete problem, including all network fees, guaranteeing that no better execution route exists.

References

[1] Fabien Guilbaud and Huyen Pham. "Optimal High Frequency Trading with limit and market orders". In: *CoRR* abs/1106.5040 (2011). arXiv: 1106.5040. URL: http://arxiv.org/abs/1106.5040.

[2] Yu Zhang, Yafei Li, and Claudio Tessone. "A Line Graph-Based Framework for Identifying Optimal Routing Paths in Decentralized Exchanges". In: *arXiv preprint arXiv:2504.15809* (2025).

[3] Guillermo Angeris and Tarun Chitra. "Improved Price Oracles: Constant Function Market Makers". In: *AFT '20: 2nd ACM Conference on Advances in Financial Technologies, New York, NY, USA, October 21-23, 2020.* ACM, 2020, pp. 80–91. DOI: 10.1145/3419614.3423251. URL: https://doi.org/10.1145/3419614.3423251.

[4] Vincent Danos, Hamza El Khalloufi, and Julien Prat. "Global Order Routing on Exchange Networks". In: *Financial Cryptography and Data Security. FC 2021 International Workshops - CoDecFin, DeFi, VOTING, and WTSC, Virtual Event, March 5, 2021, Revised Selected Papers*. Ed. by Matthew Bernhard et al. Vol. 12676. Lecture Notes in Computer Science. Springer, 2021, pp. 207–226. DOI: 10.1007/978-3-662-63958-0_19. URL: https://doi.org/10.1007/978-3-662-63958-0%5C_19.

[5] Guillermo Angeris et al. "Optimal Routing for Constant Function Market Makers". In: *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*. Ed. by David M. Pennock, Ilya Segal, and Sven Seuken. ACM, 2022, pp. 115–128. DOI: 10.1145/3490486.3538336. URL: https://doi.org/10.1145/3490486.3538336.

[6] Theo Diamandis et al. "An Efficient Algorithm for Optimal Routing Through Constant Function Market Makers". In: *Financial Cryptography and Data Security - 27th International Conference, FC 2023, Bol, Brač, Croatia, May 1-5, 2023, Revised Selected Papers, Part II*. Ed. by Foteini Baldimtsi and Christian Cachin. Vol. 13951. Lecture Notes in Computer Science. Springer, 2023, pp. 128–145. DOI: 10.1007/978-3-031-47751-5_8. URL: https://doi.org/10.1007/978-3-031-47751-5%5C_8.

[7] Stefan Loesch and Mark Bentley Richardson. "Marginal Price Optimization". In: *arXiv preprint arXiv:2502.08258* (2025).