Hidden Centralization in Decentralized Finance: Network Modeling of Liquidity Dynamics in DeFi

Nicholas Strohmeyer*†, Michael D. Norman, Ph.D.*, Yiannis G. Karavas*,
Rudra Raval†, and Sriram Vishwanath, Ph.D.*†
*MITRE
†University of Texas at Austin

***Abstract*—The decentralized finance (DeFi) ecosystem represents a complex adaptive system where liquidity providers act as critical non-deterministic agents whose behaviors collectively shape market dynamics. This whitepaper presents an integrated analysis and the fundamental building blocks for a network modeling approach to understanding liquidity provision dynamics, examining both established flow patterns and the critical relationship between liquidity removal events and price volatility. We find that tracing and aggregating the on-chain actions of individual liquidity providers (LPs) who are active at multiple pools provides a reliable signal for quantifying real-time relationships between liquidity pools and across exchanges. We establish that these *mutual* LPs are fairly widespread in historical on-chain data and demonstrate that several liquidity pools were often majority owned by just a small percentage of individual LP addresses during some of the largest periods of market volatility in 2021. Finally, we investigate and discuss some possible motivations for LP behavior. This paper represents a critical first step towards identifying and developing mitigation strategies for systemic risk in the crypto ecosystem.**

***Index Terms*—cryptocurrency, liquidity, complexity science, network science**

I. Introduction

Decentralized finance (DeFi) is a financial system based on peer-to-peer transactions built on blockchain technology. In effect, it is a complex dynamical system that emerges from the interactions of investors and traders interacting through smart contracts that encode system rules and execute transactions. The system features numerous, independently traded, and often volatile digital assets whose values may be coupled with or fundamentally derived from others. However, such relationships between digital assets and markets can be hidden, difficult to detect and may only become apparent after significant events at which point it can be too late to mitigate against risk of unwanted exposures. The permissionless nature of the complex ecosystem means that multiple on-chain markets for the same assets may emerge, effectively fragmenting liquidity.

This motivates the growing need for developing models and simulations to aid in understanding of DeFi's complex dynamics, not only for developing strategies of interaction with this complex system, but also to understand the sensitivity of the system to the influence of well-capitalized (and potentially malicious) actors. We suspect that, despite its decentralized nature, large players may have a disproportionate influence on market dynamics.

Interpreting DeFi as a networked system can be beneficial for several reasons. First, it is a natural model to capture the interactions between agents (either human or AI) and smart contracts, both in a *logical* sense, where relationships can be defined based on the rapid propagation of price dynamics (often due to movement of liquidity within the context of a single chain) and relative incentives among platforms/markets, and in an *infrastructural* sense, where various coupled blockchains (e.g., Arbitrum and Ethereum) enable the movement of liquidity between chains, sometimes with temporal delays. Similar network-based analysis of DeFi has been performed by tracing sequences of underlying smart contract calls within a chain [1]. The present work is differentiated by our focus on the economic dynamics of Liquidity Providers within the system.

The second advantage of the model is that it enables the use of existing tools for network analysis. In this report, we focus on a *logical relationship* network in which various decentralized exchanges (DEXes) are connected by the mutual volume of assets that move between them, and we present some preliminary analysis from using this approach. It should be noted that a set of *infrastructure relationships* would likely need to be integrated into the network model to meaningfully model cross-chain liquidity dynamics, especially in the context of optimistic rollups whose canonical bridge contracts bridge assets from Layer 2s to Ethereum Mainnet with week-long fraud proof delays. Although not the focus of this work, it is critical for future systemic resilience efforts to note that urgently needed liquidity traversing such a bridge would likely not arrive in time to help alleviate the financial stress of a liquidity-induced crisis.

Herein, we make the following contributions:

1) We present the concept of a *Liquidity Network* for modeling DeFi dynamics which uses the activity of *liquidity providers* (LPs) and token quantities to measure the movement of liquidity through the broader system. We present initial steps towards constructing a prototype model (provide a tangible framework for how it can be constructed and scaled) and provide our next steps towards scaling and using this model.
2) We present findings on the prevalence of *mutual providers* and their relative size, which represent an important subclass of active LPs which are observed frequently both at individual exchanges and in transferring assets between exchanges.
3) We provide an analysis of the impact LP activity at Uniswap v2-like DEXes can have on the broader USD

Email: me@mdnorman.com

Coin (USDC)/ Wrapped ETH (WETH) market. Often, active mutual providers own a significant portion of liquidity or LP share at individual points, representing potential hidden centralization among DEXes. We show that these wallets not only offer a definition of logical interconnectivity between exchanges but also that their real-time activity can potentially be used as general market indicators.

II. Background

A. Decentralized Exchanges and Liquidity Provision

Our analysis focuses on the movement of assets between *decentralized exchanges* (DEXs) [2] and the actions of *liquidity providers* (LPs) who provide assets to these exchanges. Centralized exchanges (CEXs), like Coinbase and Binance, use a central limit-order book model (CLOB) and are required to fulfill transactions off-chain as CLOB's are currently untenable for truly decentralized blockchains. In contrast, DEXs are able to make large sums of liquidity readily available to users to swap/exchange digital assets without an intermediary. Thus, DEXs have been a fundamental pillar of DeFi since they were (first) popularized by Uniswap [3] in 2018.

DEXs operate directly on their native blockchain through the automatic execution of their composite smart contract addresses (SCAs). Users can connect to the DEX via their individually-owned wallet address(es) which are defined technically as *externally-owned accounts* (EOAs). Once connected, users can initiate a set of basic transactions: "burn", "mint", "swap", etc., which are executed by the SCAs and submitted to the blockchain. In the process, the EOA address will be recorded in the block and forever associated with the transaction.

Liquidity Pools: DEXs handle swaps by using *liquidity pools*. These are smart contracts which offer a means for individuals to make liquidity available to others by collectively locking their digital assets in the pool where traders can then go to make a swap for a small fee. In exchange, these LPs are given a pro-rata share of the revenue generated from these trading fees based on their relative portion of the overall liquidity provided. In the simplest case, the LP deposits a pair of assets (WETH/USDC for instance) in equal value proportion. So, for instance, if Token A is twice as valuable as Token B, the LP is required to deposit twice the quantity of B than A. The LP is given an "LP token" which is a record of the deposit and is good for redeeming liquidity for the underlying asset pair at any future time.

Constant Function Market Makers: For liquidity pools to function, they require an Automated Market Maker (AMM) mechanism [4] to continuously set the price of one token denominated in the other(s). In early DEXs, such as Uniswap [3], the AMM does this with a simple *constant function* implementation, in which the product of the token quantities A and B before a swap is executed equals the product of the quantities after. Mathematically speaking:

$$x_t y_t = x_{t+1} y_{t+1} = k \tag{1}$$

where k is a constant. The price of one token denominated thus moves according to eq. (1) with every swap.

Two important concepts are (price) slippage and impermanent loss (IL). Slippage refers to the fact that under the constant product model, no swap ever executes perfectly at spot. However, this will be negligible if the swap is small relative to the depth of the pool. IL occurs since the value of one token relative to the other constantly moves with respect to eq. (1), the LP can only ever recover the original value of their deposit (not including trading fees accrued) if they redeem when the ratio of quantities A and B is equal to what it was when they made the deposit. This is also referred to as "divergence loss" in some literature (e.g. see [2]).

B. Liquidity Provision as a Traditional Finance Analogy

1) Options Writing: In order to better understand the risk exposure that LPs are taking, and perhaps begin to shed light on the motivations for their behavior as critically-important agents in this ecosystem, it is helpful to draw an analogy to the TradFi markets. To that end, the TradFi market participant risk exposure that most closely resembles LPing in DeFi is options writing. In the simplest strategic options contexts, i.e., selling naked calls and puts, an options writer is most profitable when the asset experiences large amounts of ranged volatility. This analogy also applies to commodities futures markets, where futures contract writers take on similar tail risk exposure to equities option contract writers.

This is because when writing options, the implied volatility (IV) is a core driver of the premium, and results in large amounts of theta, a parameter in the most commonly-used model to value options, Black-Scholes [5]. Simply put, theta determines the amount of value that the option loses over time. An options-writer exposes themselves to tail-risk for the duration of their open position, so some might consider a strategy where they close their position (by purchasing the option they sold back at a lower price) after collecting theta-derived profits.

When an LP provides liquidity to a pool, they are collecting fees on each swap done against the pool. Thus, just like someone selling naked calls or puts, they are maximally profitable when there is lots of ranged volatility. And just like an options writer, they are subject to significant losses when there is large amounts of strongly directional volatility.

Assuming a simple Uniswap v2-like AMM, or even a more sophisticated one like Curve's stableswap [6], an LP will in effect be selling the asset that experiences gains for the asset that loses value for the entire duration of the directional move. Only the most recent swap at the current price saw the LP convert the now relatively more valuable asset at the current price. Thus, once a strong directional move has occurred, and it maintains stability afterwards at the newly-balanced quantitative relationship between the two assets in the pool, the combined value of the LPs assets will be less than when they first deposited the assets into the pool. If the swap fee-based yields are also somewhat low, which normative competitive forces tend to dictate, then the LPs total position if they were to remove their liquidity at this time will have suffered IL. This is directly analogous to an options writer who experience a tail-risk event and sees the options they sold gain in value, making it more expensive to close their position (which they must do). One complication introduced by the analogy that can be disregarded is the expiration of the option. If an LP believes that their assets' relative values will strongly revert to their long-term mean, then they can avoid incurring the losses that are equivalent to closing an underwater naked options sale by simply waiting.

2) Market Making Dynamics: It's important to understand the effects of liquidity on markets. In a TradFi context, liquidity can be thought of as the order book depth. A deep order book enables lots of volume without huge amounts of price volatility via slippage. In TradFi, market makers (and traders more generally) place bids and offers in a CLOB in order to provide liquidity; there is no distinct LP agent needed. In this way, they are able to carefully define their risk exposure.

As discussed in the previous subsection, LPs make DeFi markets by effectively writing options against their holdings when depositing asset pairs into a pool. It would make sense that, much like an options writer would, they might employ a strategy to avoid strong directional volatility in order to maximize profit and avoid losses. Assuming this is the case, then a simple strategy might involve defining an expected price range to provide liquidity within and removing liquidity as soon as that range is exited.

The effect of this agent behavior is that strong directional volatility may create shallower liquidity, thereby increasing the volatility. This increased volatility then may shake other LPs out of the pool, which decreases liquidity even further. This hypothesized reflexive feedback loop may help explain the uniquely volatile price action in crypto. If so, the key to more approachable markets may be in augmenting AMMs to incentivize sticky liquidity, or perhaps even re-imagining the on-chain mechanisms for asset swapping entirely. If decentralized computation gas and blockspace were not scarce assets, a CLOB would likely be available on-chain already. Relatively centralized chains which provide high bandwidth, low latency performance, such as Solana, have attempted to implement on-chain CLOBs [7], but large-scale adoption has not taken place due to the required centralization trade-offs eroding the unique value proposition of public blockchain [8], converging on something more akin to a centralized database than a construct which produces credibly-neutral emergent trust. Given that the scarcity of these resources underpins the fundamental value proposition of public blockchain, these features are unlikely to change in the near future, and those seeking a CLOB-like user experience are likely to use a traditional off-chain CEX such as Coinbase or on-chain perpetual futures CEX such as Hyperliquid [9].

The analysis in this paper, and the analogy described above are both derived from the Uniswap v2 architecture. In v3+, the Uniswap protocol introduced the concept of 'concentrated liquidity.' This enables LPs to define the ranges at which their liquidity becomes active, enabling them to remove it per the impermanent loss avoidance strategy described above without real-time manual intervention. The impacts on strategic behavior and its effect on volatility and available liquidity, however, are the same. The advantage for those looking to model systemic risk of these v3+ architectures across the ecosystem, is that liquidity removal can be accurately forecast, opening the door for predictive and high-fidelity agent-based modeling.

III. Related Works/Literature Review

DeFi as a Networked System: A small body of previous work has also investigated DeFi from the perspective of network and complex systems science [1], [10], [11], [12]. For example, one study [1] constructed a graph of abstract protocol relationships by aggregating the frequency of underlying smart contract calls. This represents a kind of intra-chain *infrastructural approach* which tells a related and complementary story to that of a liquidity network.

Network models in crypto have also been applied to identify illicit or anomalous transactions [13], [10], [14]. One study improves performance in identifying illicit transactions using node-embedding techniques on a transaction graph [13]. The authors also found that using domain-specific features (such as block number or time between transactions) helped boost this performance over naive application of a technique like Node2Vec. This implies that out-of-the-box algorithms combined with the right feature space (perhaps mutual liquidity volumes, for instance) can be a powerful combination for building unsupervised learning models. Other notable examples use topological information to perform anomaly detection directly on Ethereum transaction graphs [10], [11].

Risk Propagation Recently, close analysis of Compound's lending protocol found that COMP rewards can negatively impact Maker's DAI [15] and it's functioning as a stable asset [16]. The authors show that COMP rewards for staking DAI put slight upward pressure on the stable price.

Furthermore, [16] also provides evidence that assets are fairly centralized in lending/borrowing protocols and that borrowers demonstrate high (perhaps unwarranted) levels of confidence in the stability of DAI, when in reality its price can be volatile. The motivational takeaway here is that relatively small deviations could lead to large numbers of liquidations, posing latent systemic risk to the wider ecosystem.

Risk can also exist at the infrastructural level. For example, evidence in [1] shows that cross-functional protocol calls at the smart-contract level are fairly common. That is, the underlying contracts of lending/borrowing protocols, for instance, commonly invoke DEX contracts to perform some function. They also conduct an analysis demonstrating high levels of direct and indirect dependence on USDT across numerous smart contracts on the largest protocols, suggesting USDT price instability could have cascading infrastructural implications across DeFi.

The logical relationship between transactions on DEXs on a blockchain (i.e., Ethereum) need not be independent of DEXs on other blockchains or CEXs either. In fact, recent work [17] claims to have found that more than 1/4th of all transactions on Ethereum's 5 top exchanges are related to arbitrage across DEXs and major CEXs. Furthermore, they estimate that about only 11 agents are responsible for 80 % of volume related this particular strategy, revealing that centralization can emerge as key drivers of broader dynamics.

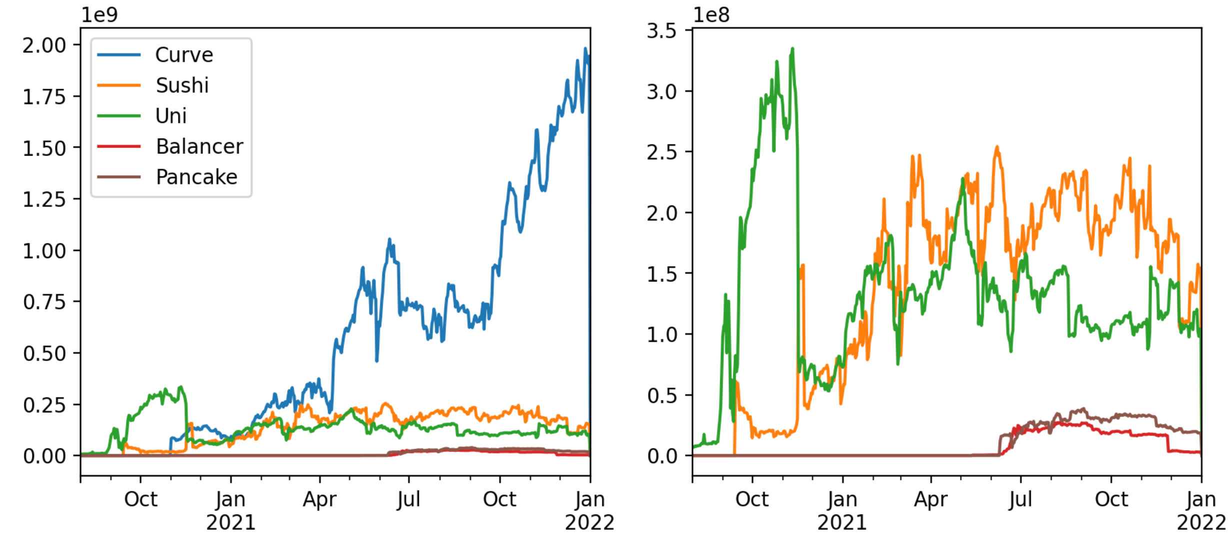

Fig. 1. **Left panel:** A history of USDC volume staked at five major v2 liquidity pools. **Right panel:** Curve is removed to show trends at lower scale.

Liquidity Provider Behavior: Similar to this work, [12] also takes an interest in the concept of a *liquidity flow* and identifies the prevalence of "overlapping" LPs across top liquidity pools on Uniswap. However, in that work, their analysis is not focused on building a networked model of liquidity provision at scale.

Furthermore, the attraction of LPs to protocol-offered incentives, a phenomenon often referred to as "liquidity mining," was the focus of recent work which specifically used the Sushiswap vampire attack as a detailed case study of LP behavior and how it affects movement of liquidity [18].

Relationship to our Objective: In summary, these previous works share various insights on themes of networked models, cascading risk, liquidity dynamics, and agent behavior. Our interest is using historical data of individual LP EOAs across points in the system as a means to measure a logical connection between them. We seek to understand the robustness of this measure for building a scalable network model to capture the dynamics of risk and liquidity propagation in DeFi and beyond.

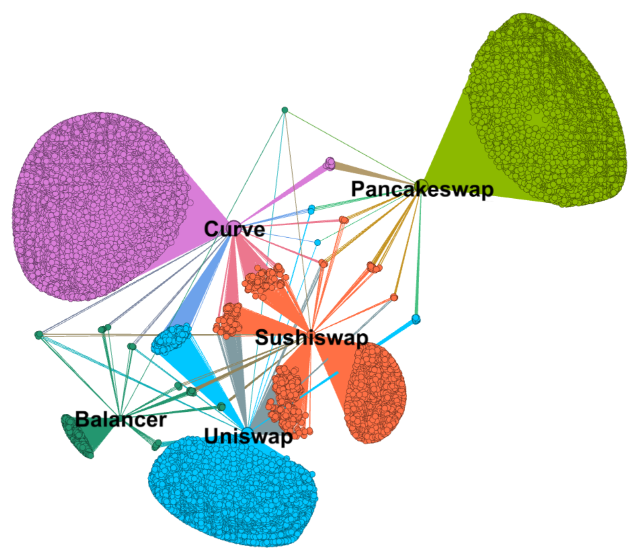

Fig. 2. Compare to fig. 3. Similar to findings in [12] a larger quantity of EOAs are observed at only a single exchange (in this set of 5). Previous work suggested many LP's are passive and make up relatively small positions [18], while a relatively small (%-wise) quantity of EOAs are the most active across mutual exchanges. As our analysis subsequently shows, some of the largest positions tend to coincide with these same EOAs.

IV. Research Framework and Methodology

As a first step towards this broader-scale DeFi liquidity model, we focus our attention on liquidity provider activity. In particular, we investigate the intersection of features which characterize liquidity providers (LPs) using liquidity pool state and event data, across 5 major decentralized exchanges during 2020 and 2021 (fig. 1). The curation of these datasets, as well as the definition of the network and key metrics enables the preliminary analyses of the proceeding sections.

A. Definitions

Mutual Provider or LP: A single EOA which provides assets to exchange A and exchange B at a given block height would be a *mutual provider*. A relaxed definition would not require a single address to *simultaneously* provide liquidity to A and B, but rather within a relatively small *time window*. In this work, we focus more on dynamic event data and typically the relaxed definition applies.

Liquidity Flow: Our aim is to measure a transfer of assets from A to B or vice versa. One possible way to measure this is to check for a net outflow at A (more LP token burn volume than mint volume) while measuring an inflow for B (excess mint volume) within a reasonable interval of time. However, when observing transactions in isolation, it can be unclear when a transfer was likely to have occurred. If instead we measure this series of transactions as having gone through the intermediary of a mutual provider address, we increase

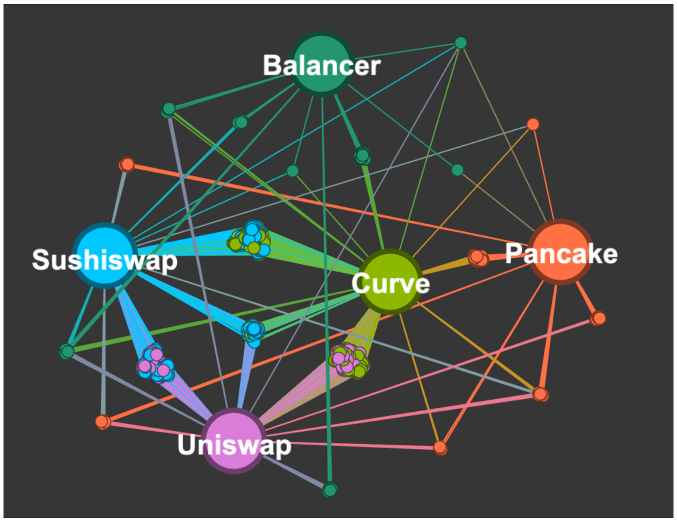

Fig. 3. Network view of the aggregate set of EOAs with observed liquidity addition or removal events for Curve's 3Pool and the other four DEXs' WETH/USDC pools from 10/2020 to 12/2021. An edge is formed between an EOA address and a DEX if that EOA burned/minted LP tokens during this period.

our confidence that these transactions do, in fact, represent a direct transfer. Thus, we define *liquidity flow* as an observation of a specific EOA removing liquidity from A (or B) and later adding at B (or A) within a configurable time window (where we use 3 days).

B. Dataset Collection

Token Balances (The Graph) One dataset was collected to measure the static state of mutual liquidity provision between Uniswap and Sushiswap. We called the *liquidityBalance* contracts on multiple sub-graphs (one for each exchange) from The Graph protocol. This data provides user EOA addresses and the amount of LP tokens held by the particular address at a given block height, allowing one to construct the complete unique set of liquidity providers to a particular exchange at a particular time by aggregating over non-zero balances.

Liquidity Events We chose to focus primarily on analyzing dynamic relationships for which we built a second dataset from *liquidityEvents* for the 5 exchanges listed in table I from September 2020 through December 31, 2021. Each time-stamped event is specific to a liquidity pool (searchable by unique pool address) and includes the token symbols and amounts included in the transaction, the block in which the transaction was included, whether the event represented a mint or a burn, and, crucially, the unique EOA address of the LP who initiated the event. This last attribute allows us not only to use our definition of *liquidity flow* above but also to construct individual EOA histories across all exchanges.

Constructing individual EOA histories enables insights into general trends in LP behavior. For instance, in this work we use it to inform several case studies. In future work, we might consider using individual EOA traces for classification tasks such as clustering LPs by size/strategy patterns. It should be noted that the centralization of liquidity we uncover in this work (i.e., the prevalence and significance of mutual

TABLE I

Liquidity Pool	**Contract Address**
Uniswap v2 (USDC/WETH)	0xb4e16d0168e52d35cacd2c6185b44281ec28c9dc
Sushiswap (USDC/WETH)	0x397ff1542f962076d0bfe58ea045ffa2d347aca0
Balancer (USDC/WETH)	0x96646936b91d6b9d7d0c47c496afbf3d6ec7b6f8
Curve 3Pool (USDC/USDT/DAI)	0xbebc44782c7db0a1a60cb6fe97d0b483032ff1c7
Pancakeswap (USDC/WETH)	0xea26b78255df2bbc31c1ebf60010d78670185bd0

providers) represents a *best-case* scenario, as decentralized liquidity sources are a passive risk-reducing feature of this ecosystem; there is nothing preventing an LP from using multiple EOAs, which would make actual liquidity centralization more severe than it appears on-chain. The reason that decentralized liquidity sources reduce risk is because the subsequent slippage impact on other ecosystem participants of an agent's decision to remove liquidity is reduced if they are not individually providing significant amounts of liquidity. Thus, given the assumption that decentralized control inherently leads to decorrelated behavior, the more decentralized the liquidity ownership is, the better.

Pool Reserves (State) In addition to events, we query daily token balances at the 5 DEXs for the duration of the time period. This allows exact measurements of token quantities over time, giving us the 'state' of the pool (the node) itself.

ETH Market Data We take a complete history of Ethereum hourly market spot prices and hourly volume from an online source [19] which is used for several analyses in comparing mint/burn events to broader volume and price volatility metrics.

C. Volatility Metrics and Analysis

To measure market volatility we examined (i) the day-over-day difference in price and (ii) the standard deviation of hourly log-returns over a 30-hour rolling window. We found these two metrics were often highly correlated and therefore only present the former for its simpler interpretation.

D. Network Construction

Node Definition: A node in the liquidity network represents a frequently used smart contract relevant to DeFi operations. In the scope of this report, nodes are smart contracts associated with DEXs – in particular WETH/USDC liquidity pools and several related markets on PancakeSwap [20], Balancer [2], and Curve's 3pool [21]. Fundamentally, we are interested in the relationship between these DEXs and their users. To that end, we may choose to represent significant individual user EOAs as additional nodes in the system. Given that they are all ontological peers, i.e., Ethereum Virtual Machine (EVM) addresses capable of participating in transactions, this makes intuitive sense.

Edge Definition In fig. 2, nodes include *all* liquidity providers and edges represent their liquidity provision relationship to a pool. However, in fig. 3, isolated LP's are dropped and only edges incident through *mutual providers* remain. Thus, we can measure the strength of a connection between two DEXs or, more generally, two smart contracts, using the

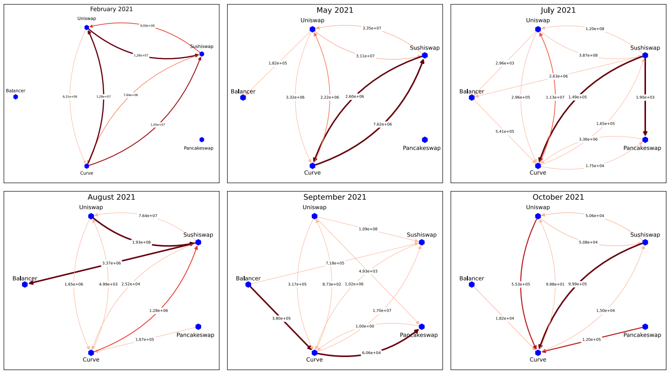

Fig. 4. A dynamic representation of fig. 3 chronologically arranged (left to right) for select months in 2021. Volume transferred between exchanges (as measured through mutual LP's) is aggregated over a 3-day sliding window. The edge thickness depicts volume between particular exchanges which varies with time. Edge directions depict net movements.

volume transferred over these links through mutual wallets as depicted in fig. 3, which is how we measure overall shared liquidity between the two exchanges.

Dynamic Network Connections between nodes will form and vanish over time making the network fundamentally dynamic. Furthermore, individual nodes may grow or fade in relevance (liquidity, frequency of activity, etc.). In fact, prior work suggests that a large group of LPs in the DeFi environment of 2020/2021 represent *passive* providers. If a DEX were to contain a high percentage of these passive participants, its relationship to other nodes may fade completely. Thus, a liquidity dynamics-based network model (illustrated in fig. 4) would constantly update the strength of connections and the state of the nodes based on recent and current observations within the aforementioned 3-day sliding window interval. Note that this interval is a configurable parameter. If the window is too short, it may miss out on logically-related events; if it is too long, it will be overly inclusive of spaced-out events that may have little-to-no relation. We found that 3 days represented a fair compromise between these extremes.

V. Network Structure Analysis

In this section, we present our analysis into the prevalence of "mutual LP" connections across popular decentralized exchanges from September 2020 through the end of 2021. We start with a case study of the logical connection between Uniswap v2 and Sushiswap WETH/USDC during the 2020 vampire attack and show that this period coincided with high levels of measurable mutual LP activity. We then expand the scope of this methodology to WETH/USDC pools across 4 different exchanges as well as Curve's stable-pool as shown in fig. 1, and then across different asset pairings within Sushiswap. While the Uniswap/Sushiswap connection in late 2020 is an illustrative extreme case, the latter two results suggests this approach is applicable to the modeling of DeFi liquidity more generally.

Uniswap/Sushiswap Connection: Both fig. 5 and fig. 6 demonstrate the "static" mutual share between Uniswap and Sushiswap WETH/USDC pools using data curated from The Graph protocol. This starting point verifies the prevalence of mutual LP's. We note that academic validation of our first-principles thinking which led to this approach was discovered via conceptual convergence [12]; we were unaware of this related work at the time of conducting our own analysis but became cognizant during the compilation of our results.

Vampire Attack: The underlying reason for the strength of the connection between Sushiswap and Uniswap in the fall of 2020 was Sushiswap's vampire attack on Uniswap. For an in-depth explanation, see [18]. It will suffice to understand just a few key facts: (i) Sushiswap offered rewards for LPs to stake their Uniswap LP tokens to a Sushiswap-controlled smart contract, then (ii) used these tokens to directly siphon liquidity away from Uniswap to Sushiswap, and (iii) in response to this, Uniswap launched their own token, UNI, thereby sparking a period of competition (essentially a bidding war for liquidity). In [18], they carefully document the fact that LPs were actively swapping between the DEXs during this time. However, the relationship persists well into 2021 (fig. 1).

Mutual Provider Relationships on 5 Select Exchanges Extending the search for mutual EOAs to the 5 exchanges listed in table I, we find that, in fact, each of these exchanges shared at least a few mutual providers throughout the period under consideration. Perhaps not surprisingly, these connections are particularly strong between Curve, Sushiswap, and Uniswap – which were and continue to be some of the deepest pools in all of DeFi. Interestingly, although Curve's 3Pool is not directly connected to the WETH market, it is the most densely connected exchange of all 5. This is likely representative of the fact that LPs will move into stablecoins as safe haven assets, and Curve is the deepest on-chain liquidity source for stablecoins. More discussion of Curve is included later. A visualization of all exchanges and their liquidity providers is given in fig. 2 whereas fig. 3 reduces this picture to mutual providers only.

Broader System Connections Next, we show that broadening the search to include *all* top liquidity pools across Sushiswap (beyond the WETH/USDC pairing) while narrowing the window of mutual activity to a 3 day period still results in a high amount of mutual activity. This finding indicates that mutual LP activity, which we spend the remainder of this paper investigating, is a highly-prevalent, robust system signal for a scalable model. Visualization of this result can be found in fig. 18

VI. LP Market Impacts and Motives Analysis

A. Timeline of Market Impacts

In this section, we highlight a few notable events within the period of September 2020 through December 2021 in which LP activity on Sushiswap and Uniswap (in particular) temporally correlate with general market volatility. On several occasions, we find that significant ETH price movements may have been directly caused by few (or even single) large LP burn events. Also of significance is the fact that the EOAs which initiated the events often fall under the dynamic definition of a *mutual liquidity provider*. Finally, we find some evidence which backs the idea that singular large LP

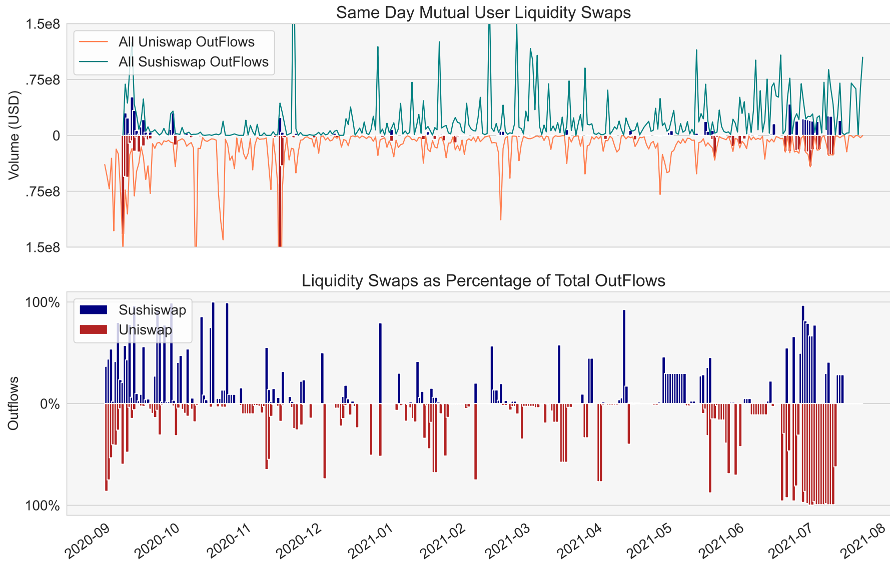

Fig. 5. **Top:** Aggregated outflows (burns) from Uniswap and Sushiswap overlaid with outflows accounted for by detectable same-day swaps (the same EOA burns at one pool and mints at the other within a 24 hour window). **Bottom:** the same data but as a percentage over time. The height of the bars is that day's % outflow accounted for by mutual users. These results suggest a nontrivial amount of daily volume at these liquidity pools is accounted for by short-term mutual activity.

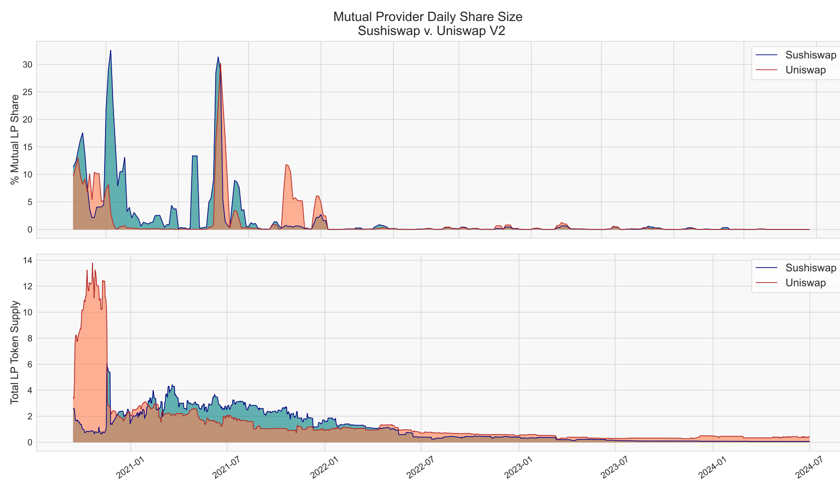

Fig. 6. **Top:** Mutual user share between Sushiswap and Uniswap WETH/USDC v2 pools over time. Specifically, there is a nontrivial set of liquidity providers holding assets on both exchanges simultaneously; a relationship that can be identified just through cross-referencing addresses. In other words, this is an effective minimum overlap, whereas actual percentages of overlap may be much higher. **Bottom:** Total LP token supply for both pools which shows diminishing liquidity over time. The period for which liquidity is highest coincides with the period in which mutual users are most staked in these pools.

events precede periods of increased mid-sized volatility events (suggesting a potential feedback loop between LP behavior and market volatility).

*1) **Tracing a November 2020 Liquidity Flow**:* Perhaps the most interesting event is in November 2020 (fig. 7), coinciding with the fallout of Sushiswap's vampire attack on Uniswap. A single LP moves assets from Uniswap to Sushiswap and eventually burns. Several days later, there is a spike in market volume and a significant downturn in price. On the one hand, this is an example of the impact of hidden centralization in a new and emerging market. On the other hand, it is an example of how an individual *liquidity trace* tells a larger story. This also relates to how understanding the differences in LP incentives (at least during this period) could be predictive of large liquidity flows.

In the spring of 2021 fig. 8, ETH experienced increased periods of volatility. In March, a single large spike in Sushiswap burns precedes the period of high volatility and is followed by a period of increased burns. This could suggest the type of negative feedback loop in which a shock (large burn) and a period of volatility leads to continued one-sided burning (see table II).

An important factor is the depth of the pool itself, which can be instantly affected by these single large LP events, and then

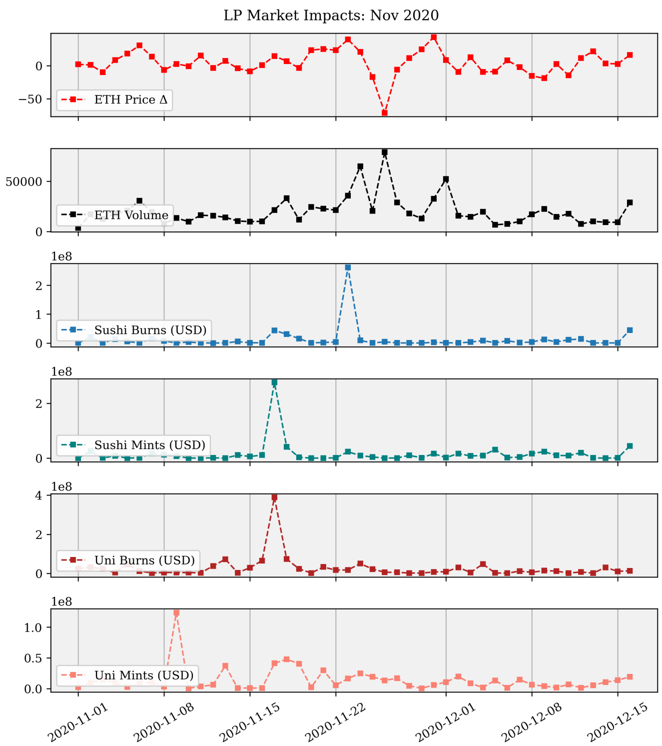

Fig. 7. Overlaying Uniswap and Sushiswap liquidity events with total market volume and day-over-day Ether price differences reveals an interesting pattern following Uniswap v2's introduction of the UNI token. The LP appears to move assets from Uniswap to Sushiswap, and then likely sells in the next few days (perhaps to cash in on gains from rewards). Interestingly this address was associated with Justin Sun (founder of Tron) using Etherscan.

potentially exacerbated by mid-sized burns. The phenomenon is shown in fig. 9.

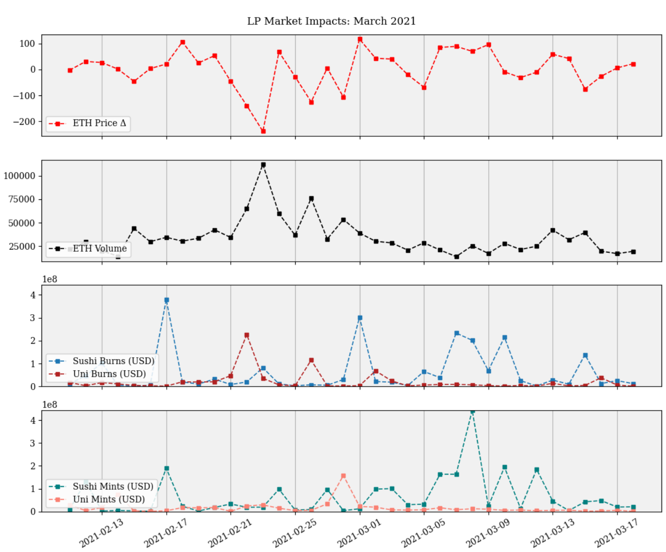

Fig. 8. Individual large wallet activity on February 17th and March 1st 2021 precede periods of high volatility in both market volume and price. Furthermore, increased burning activity in early March follows these activity spikes indicating a possible negative feedback effect. Meanwhile, the spikes are also followed by increased minting activity which tends to correlate with burns more generally.

TABLE II
EOA MOVEMENTS AROUND FEB 17, 2021 (SUSHISWAP USDC/WETH POOL

Date	Time	EOA	Volume
2021-02-17	00:57:24	0x2...f53a	1.783239e+02
	01:38:06	0xc...ff67	2.832808e+03
	05:01:51	0x7...cdd0	1.697877e+02
	07:40:52	0x7...f4f1	4.798600e+04
	08:30:56	0x6...71da	1.314644e+05
	08:33:57	0x9...5ae3	7.145202e+04
	10:36:34	**0x8...7b83**	**1.844033e+08**
	12:34:39	0x9...c4e4	1.182418e+05
	15:51:28	0xa...f807	1.690059e+04
	19:12:22	0x4...bb49	4.924204e+04
	19:29:07	0x7...2471	5.267431e+05
	22:39:16	0x3...ff1a	4.456183e+06
2021-02-18	03:54:18	0xa...c8a6	9.153682e+03
	10:31:40	0x3...ff1a	4.191405e+06
	11:06:45	0x9...56a2	7.429623e+05

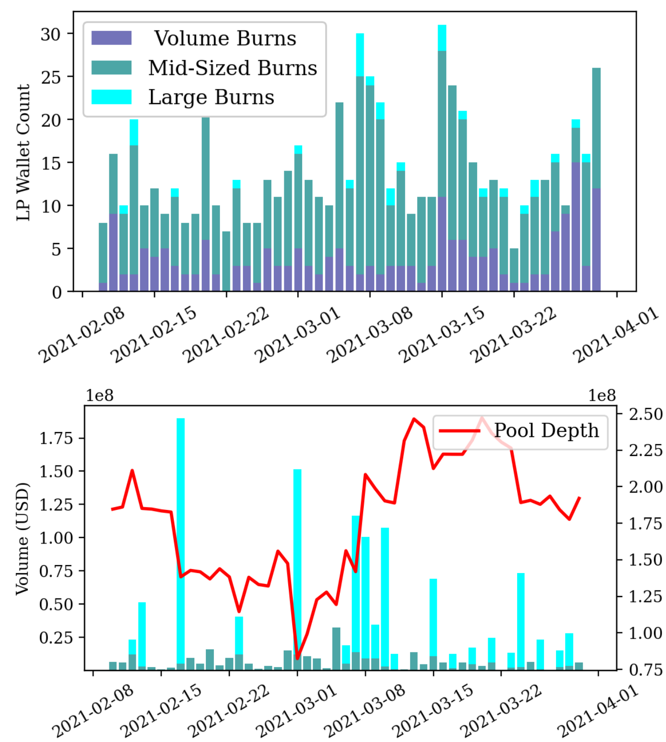

Fig. 9. LP Event Tiers. "Large" burns are defined as those which are greater than 1 million USD. Large transactions such as these are quite common. Indeed, the history of these two exchanges across 2021 is very much determined by large players such as these. Mid Sized is the next tier including transactions greater than 100,000 USD. Clearly, a small number of LPs on Sushiswap drove majority of burn volume. The large burn on March 1st correlates with a sudden spike in market volume. In an emerging market such as DEXs in 2021, large transactions such as these have much greater potential to send lasting volatility shocks through the market.

2) ***May 2021:*** May (fig. 11) experienced the largest ETH volatility of 2021. The precise cause for the steep drop in ETH demands in this case is up for debate (Tesla news). Some analysts however pointed out the moves were largely driven by "retail investors." However, the existence and correlation of these large-scale one-sided burns next to price volatility tells a slightly more nuanced story perhaps. These DEXs, which

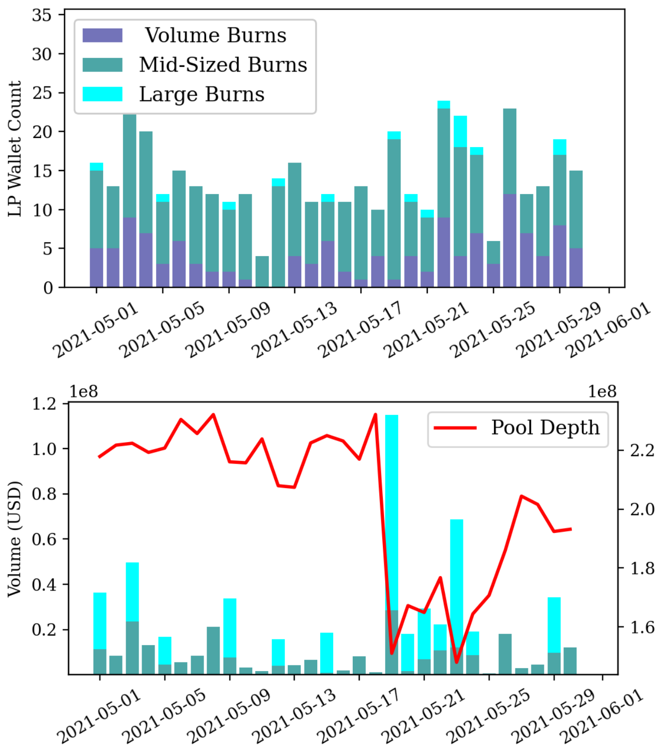

Fig. 10. LP Event Tiers. May 2021 observes a similar pattern as March 2021. Again we see just a select few EOA's having large effects on pool depth with relative minimal effect with small to mid tier events. Especially note the sharp drops in total pool liquidity which is slow to recover. On both 5/17 and 5/21, these were driven by single-EOA events.

were critical to the WETH/USDC market at this time, were largely impacted by hidden centralized capital (LPs), as shown in fig. 10.

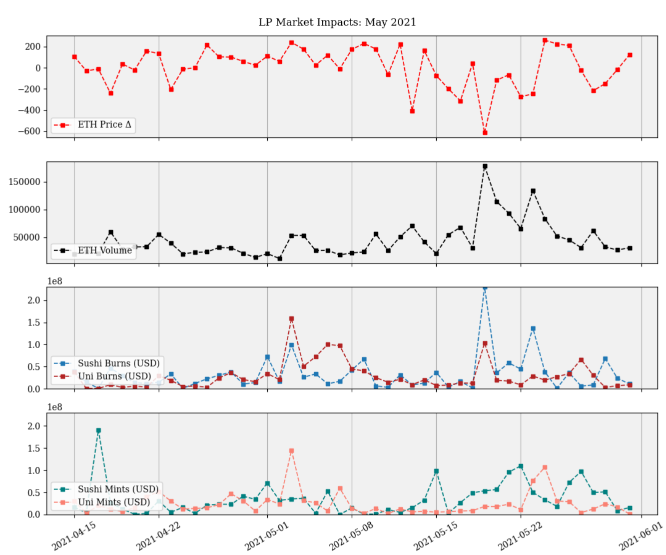

Fig. 11. High market volatility within the month likely motivates both Uniswap and Sushiswap LPs (who may also be overlapping) to pull liquidity near the end of the month, exacerbating the situation.

*3) **Fall 2021**:* Lastly, the same pattern appears again in September 2021 (fig. 12) and several spikes (where market

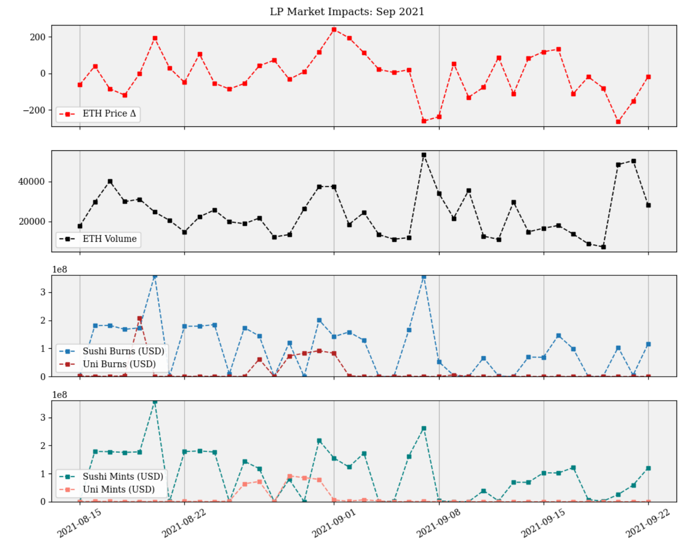

Fig. 12. September 2021, burns initiated by a single LP in particular appears to precede a spike in volume to the market, exacerbating downward price pressure. Meanwhile, downward movements in prior days may have motivated the moves. The large burn was followed by increased mid-sized burns later in the day.

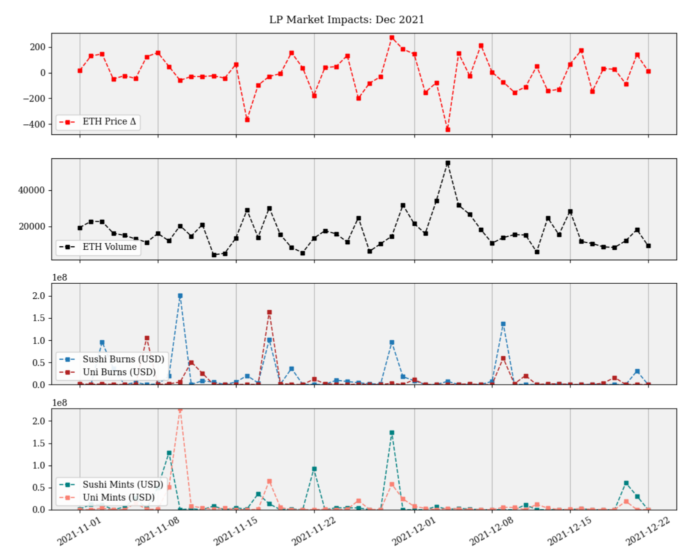

Fig. 13. November 9th, we observe a large LP transferring liquidity out of Sushiswap to Uniswap. November 16th and December 8th a mid-to-large-sized LP removes liquidity from both platforms simultaneously

impact is not as severe) on November 16th and December 8th fig. 13 are initiated by a mutual LP.

B. Other DEXs

*1) **Pancakeswap & Balancer**:* While Pancakeswap and Balancer's WETH/USDC pools represent a much smaller market than Uni/Sushi, one can observe moments when such pools might still exhibit events of significance on the broader market driven by a select few addresses, as seen in fig. 14. Furthermore, liquidity removals on smaller platforms may yet have out-sized effect on the market as shallower pools with a large enough swap volume could have volatility implications. If arbitrage liquidity is not immediately available from infrastructurally-adjacent markets, a shallower market price dislocation of an otherwise relatively liquid asset may trigger a contagion. This is analogous to the 'tail wagging the dog', where fragmentation creates pockets of shallow liquidity that may cause increased volatility due to higher slippage such that a relatively small perturbation (i.e., a swap that is insignificant when compared to the total market capitalization of an asset) can trigger nonlinear responses that propagate across the logical network faster than arbitrage can be executed across the infrastructural network to rectify what would otherwise be a localized dislocation. The impact and potential mitigation of such systemic risks may be explored in future work.

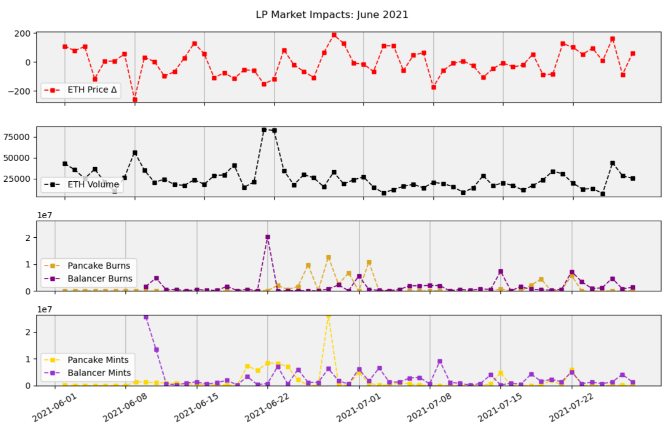

Fig. 14. The Balancer outflow spike in June was initiated by a single EOA

As a notable aside, the EOA address responsible for Balancer's large burn on June 23rd was a particularly active mutual LP. From September 2020 to the end of 2021, this address provided 1,574,205 in USDC alone in Curve's 3Pool, had over 63 events on Sushiswap, and was active during the 2020 vampire attack. It is also interesting that this LP removed liquidity on the 23rd of June, having only just added it 6 days earlier, a pattern perhaps consistent with active, opportunistic liquidity mining [18].

*2) **Curve**:* As Curve's 3Pool grows in volume it's connection to the WETH/USDC likely grows as well. It would make sense that this is where many large investors hold liquidity as a safe haven, with stablecoin yields coming from swap fees as well as CRV token incentives. A net removal from this pool may signal a preparation for a larger play in some other market, such as WETH/USDC, or vice versa. Some evidence of these dynamics are described in fig. 15. Famously, an early catalyst of the Terra Luna UST stablecoin's collapse was a large entity dumping a huge amount of UST on Curve, driving the UST-specific stablecoin pool far out of balance [21]. Utilizing these protocols for systemic ecosystem risk monitoring by focusing on centralized liquidity, not just large holders, seems to hold promise as a conceptual approach to getting further 'left of boom'.

Although numerous mutual LPs between Sushiswap and Uniswap also appear at Curve throughout 2021, the flurry of activity in May is driven mostly by EOAs distinct from those LPs. This suggests that Curve is likely a hub outside of the relatively small and tightly coupled subsystem we analyzed in

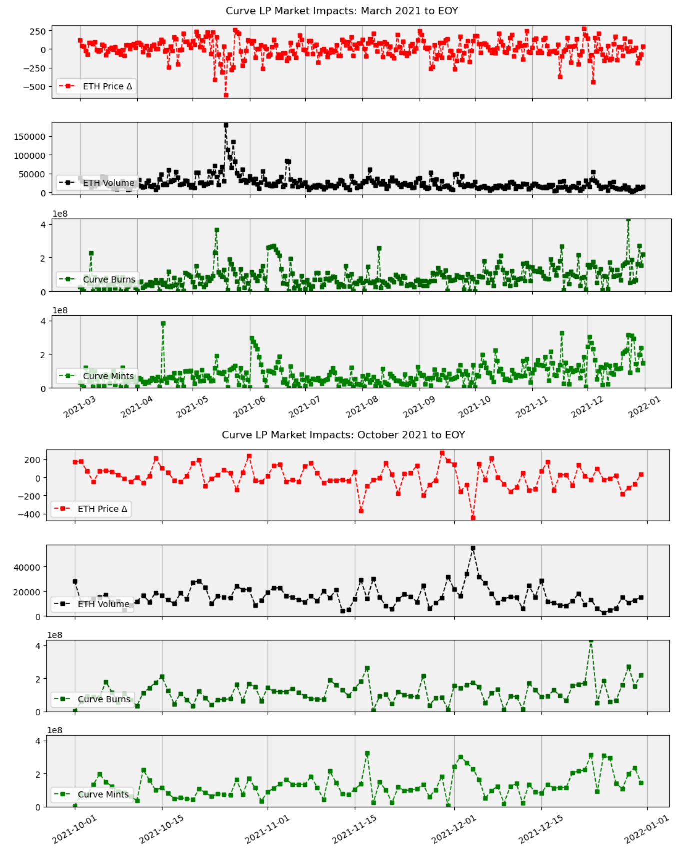

Fig. 15. **Top panel:** A large spike in burn activity on Curve's 3Pool immediately precedes some of the largest price movements of the year (March 2021). Perhaps the burning activity was motivated by the increase in velocity already building. The burns were widely distributed among many mid-to-large LPs. **Bottom Panel:** Spikes in Curve activity more broadly appear to have relation with volatile periods in Ethereum market. Net inflow of December corresponds to downward turn in ETH.

this paper. The centrality of Curve to a broader DeFi liquidity network and its growth over the past few years will likely make it an important subject of subsequent work.

C. LP Activity / Negative Feedback

In section II-B2 (Market Making) we mention that a negative feedback cycle could arise in cases where individual large burns appear to motivate continued burning from other LPs. Here, we investigate this idea by drilling further down into categories of LPs (see fig. 9). Large LP agent burns are correlated with same-day increases in medium to small burn volume which suggests there may be agent-to-agent behavioral influences at play, i.e., LPs 'copytrading' each other in the same way that smaller EOAs do to larger EOAs in the context of asset swapping.

Furthermore, single large burns such as the one shown in fig. 10 (for context, especially with respect to tokens with highly centralized liquidity provision, this type of event is referred to as a 'rug pull', as the withdrawal of significant liquidity is effectively the withdrawal of the base asset-denominated capital that has accrued to the token) make the pools shallower. With overall less liquidity in Sushiswap (in this example), the volume threshold at which swappers will experience significant slippage (i.e., impact on price) is lowered, and greater volume-driven price volatility is deterministically guaranteed to follow.

More evidence that large burns can motivate small-to-medium LP agents simply comes from the timing of events on March 1st and likewise on February 22nd. In table II, prior to the large burn, activity averages lower, and mid-sized LP activity picks up after. However, this does not immediately imply a negative feedback pattern since some burning activity is also accompanied by increased minting. Still, in this instance, the burn initiated by this large LP precedes a net outflow of liquidity. A more volatile period in both LP activity and broader ETH prices then follows.

D. Same-Day Removals and Additions

Several interesting LP behaviors are highlighted in fig. 12. A single LP EOA is the primary driver of the burn spike on both days of peak burn shown (August 20th and September 9th), and this LP exhibits a sort of "day-trading" pattern in which they burn temporarily and reappear, adding liquidity a short time later. On August 20th, they reappear after burning with a similar amount of assets. On September 9th, they removed more value than they minted which thus led to a net outflow. Also, on September 9th, they exhibited this pattern on two separate occasions – the first time within 15 minutes and the second time over multiple hours. This suggests that some of the largest LPs are running active risk mitigation strategies. While the incentive to side-step volatility is clear to individual LPs, at scale this behavior may present an emergent pathology to the market. It is paradoxical to consider that individual agent strategies to reduce risk may lead to an increase in systemic risk.

E. Ranged Volatility Behavioral Hypothesis

Recalling a hypothesis from section II-B2 (Market Making), we designed a simple analysis to test how often LP actions (burns in particular) could relate to the DEX's WETH price deviations from a short-term historical *range*. Given periods of extreme volatility, the timing of these exits may not be surprising, as we have already seen they tend to correlate with price dips. We also saw evidence in the previous section that such large burns in particular may accelerate these dips. Interestingly, from the perspective of ranged volatility, we also find evidence that some of the largest burns throughout 2021 could be *responses* to increased price volatility. In particular, if we set up a rolling 10-day 20th and 80th percentile price range (floor and ceiling, respectively) we find that exits from this range correlate with net outflows (initiated by LPs) from the given liquidity pool. These results are documented in fig. 16 (Uniswap) and fig. 17 (Sushiswap).

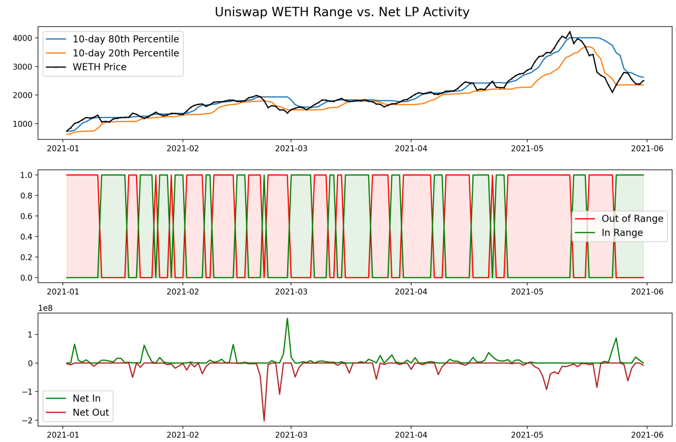

Fig. 16. Top: Uniswap WETH/USDC price series versus its 10-day rolling range. Middle: Indicators of when price is in or out of range. Bottom: Net inflow/outflow from LP activity in this period. We note that WETH price exits from range may explain the timing of many outflows. Inflows often follow outflows (either due to incentive or perhaps when price is back in range). The largest burns seem to coincide with the largest drops (when prices break the lower limit of the range). This pattern suggests a possible explanation of net mint/burning activity.

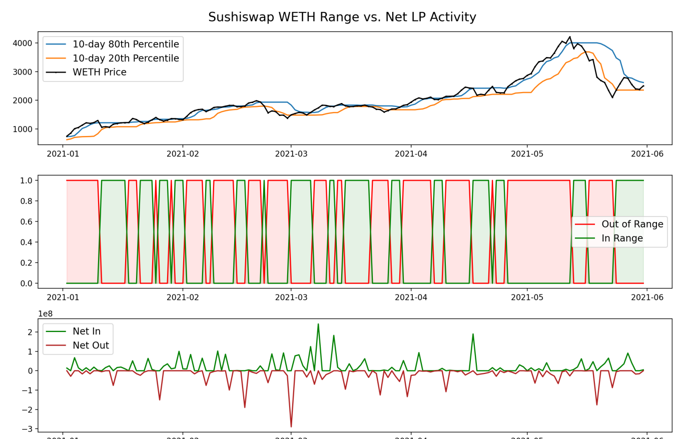

Fig. 17. The same information presented in fig. 16, but with Sushiswap 2021 data. Again many of the largest burns are timed precisely around moments when the WETH prices breaks through the 10 day rolling range pattern. LP activity spikes seem to appear near transition period in particular. A slight bias towards net outflows exists when the price breaks the 10-day floor versus exceeding than ceiling.

VII. Conclusion

In this report, we have presented the concept of a network model of liquidity dynamics. This inherently dynamic representation is aimed at capturing relationships between critical nodes in the DeFi ecosystem which, in the most general case, would include large liquidity providers, decentralized exchanges, lending/borrowing protocols, and perpetual futures markets, among others. Through the further development of this model, we hope to capture key features of the bottom-up behavioral dynamics governing decentralized finance, and even its potential relationship to the TradFi ecosystem. The scope of this report has covered several v2 liquidity pools which were critical to the WETH/USDC market in the 2020/2021 era of DeFi.

We defined the concept of *mutual liquidity providers*, unique EOA addresses that provide liquidity in multiple markets of the same asset simultaneously or within close temporal proximity, as a means of identifying hidden centralization of liquidity control, and thus latent systemic risk, should that liquidity be removed from one or more markets in size. Our analysis finds the mutual liquidity provider concept is robust for measuring logical economic relationships and tracking flows.

By extending our scope with the inclusion of cross-chain pools (i.e., Pancakeswap on Binance's BNB Chain), different AMM variants (Balancer/Curve), and token pairings beyond WETH/USDC (fig. 18 in Appendix), we find mutual LPs at a broader scale. Furthermore, many of these mutual LPs control enough liquidity that their actions have apparent impact on both local and general market price dynamics. This suggests to us that (i) tracking of individual EOA actions is scalable for modeling large DeFi dynamics (ii) these EOA actions are likely important signals for interpreting the system and (iii) that both of these points together can be used in developing simulations and predictive ML tools.

Furthermore, we observed large degrees of capital centralization in a small neighborhood of DEXs from the 2020-21 period crucial to the WETH/USDC market. In fact, our analysis considers the "best-case" scenario in which we assume a single EOA address represents a single agent. In reality, the level of centralization could be far worse than what the naive assumption would suggest since there is no limit on how many EOA addresses an individual agent can own and operate.

Motivated by the above points, we provided a preliminary analysis of factors that might drive LP behavior, especially in the mid-sized to larger events. We hypothesized that arbitrage between DEXs and CEXs plays an important role in the timing of events, although given the scale and resolution of our dataset, more analysis is needed, which we plan for future work. We also offered connections to market-making and the writing of options/futures contracts in traditional finance. From this perspective, the LP agent benefits most when the pool's WETH price lies within a predictable range such that fee revenue will outweigh the expected impermanent loss inherent to liquidity provision. We found several instances that exemplify this in which large LPs can be observed exiting the pool once the ETH price broke outside a seven-day 60% rolling confidence range.

In some cases, we saw that large burns were followed by an increased number of mid-sized LP burns, suggesting the possible existence of agent-to-agent feedback loops, indicating an agent-based simulation approach may be warranted. Paradoxically, rational behavior that limits individual LP agent risk, at scale, creates pathologically emergent systemic risk via negative impacts on slippage and the resulting volatility.

A. Future Work

Two immediate lines of work follow, including (i) scaling of the liquidity network and (ii) targeted LP agent behavioral analysis. The first involves inclusion of a broader set of node/node-types in the DeFi ecosystem as well as working to define metrics for cross-chain/Layer 2 and CEX-DEX arbitrage relationships. Also, we can begin applying ML techniques to analyze these liquidity networks. Some techniques that appear to be immediately applicable include unsupervised node embeddings like Node2Vec, with features crafted to this domain.

The second direction involves classifying LP agent behavior and the signals they most commonly respond to, allowing us to build an LP ontology for v2/v3/v4 DEXs, with replicability for other related protocols. Such work would enable agent-based simulations informed both by the real-world data component and the structural component provided by the liquidity network. Both directions have solid grounding in findings from this report as well as previous work that applied techniques like behavioral clustering [18] and node embeddings [13].

These near-term research directions will contribute to the larger goal of mitigating systemic risk in the DeFi ecosystem with behavioral monitoring and participatory response mechanisms. Finally, recommendations for improving market structure via augmentation of incentive mechanism design to include emergent property-based feedback loops which align individual LP agent risk mitigation strategies with systemic risk mitigation needs are also an envisioned outcome of future work.

As DeFi and public blockchain technology continues to grow, it is crucially important to continue developing frameworks for predictive and scenario-testing tools that would enable concerned participants to interact with this system strategically and informed by data. Our approach can provide insight on the system level (network) and the agent level (bottom-up behavioral dynamics). Such tools can offer the capability to identify and plan for the potential of misaligned actors who, with sufficient capital, can have dramatic, destabilizing impact whilst operating within the intended context of DeFi.

References

[1] S. Kitzler, F. Victor, P. Saggese, and B. Haslhofer, "Disentangling decentralized finance (DeFi) compositions," *ACM Transactions on the Web*, vol. 17, no. 2, pp. 1–26, 2023.

[2] J. Xu, K. Paruch, S. Cousaert, and Y. Feng, "Sok: Decentralized exchanges (DEX) with automated market maker (AMM) protocols," *ACM Computing Surveys*, vol. 55, no. 11, pp. 1–50, 2023.

[3] Y. C. Lo and F. Medda, "Uniswap and the emergence of the decentralized exchange," *Journal of financial market infrastructures*, vol. 10, no. 2, pp. 1–25, 2021.

[4] V. Mohan, "Automated market makers and decentralized exchanges: a DeFi primer," *Financial Innovation*, vol. 8, no. 1, p. 20, 2022.

[5] J. D. MacBeth and L. J. Merville, "An empirical examination of the black-scholes call option pricing model," *The journal of finance*, vol. 34, no. 5, pp. 1173–1186, 1979.

[6] M. Egorov and C. Finance, "Automatic market-making with dynamic peg," *Retrieved Dec 2021*, 2021.

[7] R.-A. Tănase, C. Carabaş, and N. Țăpuș, "Limit order for decentralized exchanges on blockchain," in *2023 22nd RoEduNet Conference: Networking in Education and Research (RoEduNet)*. IEEE, 2023, pp. 1–11.

[8] M. D. Norman, Y. G. Karavas, and H. Reed, "The emergence of trust and value in public blockchain networks," in *IX International Conference on Complex Systems*, 2018, p. 22.

[9] J. Scharfman, "Decentralized finance (DeFi) fraud and hacks: Part 3," in *The Cryptocurrency and Digital Asset Fraud Casebook, Volume II: DeFi, NFTs, DAOs, Meme Coins, and Other Digital Asset Hacks*. Springer, 2024, pp. 159–191.

[10] Y. Li, U. Islambekov, C. Akcora, E. Smirnova, Y. R. Gel, and M. Kantarcioglu, "Dissecting ethereum blockchain analytics: What we learn from topology and geometry of the ethereum graph?" in *Proceedings of the 2020 SIAM international conference on data mining*. SIAM, 2020, pp. 523–531.

[11] D. Ofori-Boateng, I. Dominguez, M. Kantarcioglu, C. Akcora, and Y. Gel, "Topological anomaly detection in dynamic multilayer blockchain networks. 2021."

[12] L. Heimbach, Y. Wang, and R. Wattenhofer, "Behavior of liquidity providers in decentralized exchanges.(2021)," *URL https://arxiv.org/abs/2105.13822*, 2021.

[13] F. Poursafaei, R. Rabbany, and Z. Zilic, "Sigtran: signature vectors for detecting illicit activities in blockchain transaction networks," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2021, pp. 27–39.

[14] D. Ofori-Boateng, I. S. Dominguez, C. Akcora, M. Kantarcioglu, and Y. R. Gel, "Topological anomaly detection in dynamic multilayer blockchain networks," in *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part I 21*. Springer, 2021, pp. 788–804.

[15] J. Chen, K. Fogel, and K. John, "Understanding the maker protocol," *arXiv preprint arXiv:2210.16899*, 2022.

[16] D. Perez, S. M. Werner, J. Xu, and B. Livshits, "Liquidations: DeFi on a knife-edge," in *Financial Cryptography and Data Security: 25th International Conference, FC 2021, Virtual Event, March 1–5, 2021, Revised Selected Papers, Part II 25*. Springer, 2021, pp. 457–476.

[17] L. Heimbach, V. Pahari, and E. Schertenleib, "Non-atomic arbitrage in decentralized finance," *arXiv preprint arXiv:2401.01622*, 2024.

[18] S. Fan, T. Min, X. Wu, and C. Wei, "Towards understanding governance tokens in liquidity mining: A case study of decentralized exchanges," *World Wide Web*, vol. 26, no. 3, pp. 1181–1200, 2023.

[19] investing.com, "Ethereum (ETH/USD)." [Online]. Available: https://www.investing.com/crypto/ethereum/historical-data

[20] O. Mazor and O. Rottenstreich, "An empirical study of cross-chain arbitrage in decentralized exchanges," in *2024 16th International Conference on COMmunication Systems & NETworkS (COMSNETS)*. IEEE, 2024, pp. 488–496.

[21] A. Briola, D. Vidal-Tomás, Y. Wang, and T. Aste, "Anatomy of a stablecoin's failure: The terra-luna case," *Finance Research Letters*, vol. 51, p. 103358, 2023.

Appendix

Mutual Liquidity Adjacency Matrix for Top Asset Pairings on Sushiswap (Sep - Dec 2020)

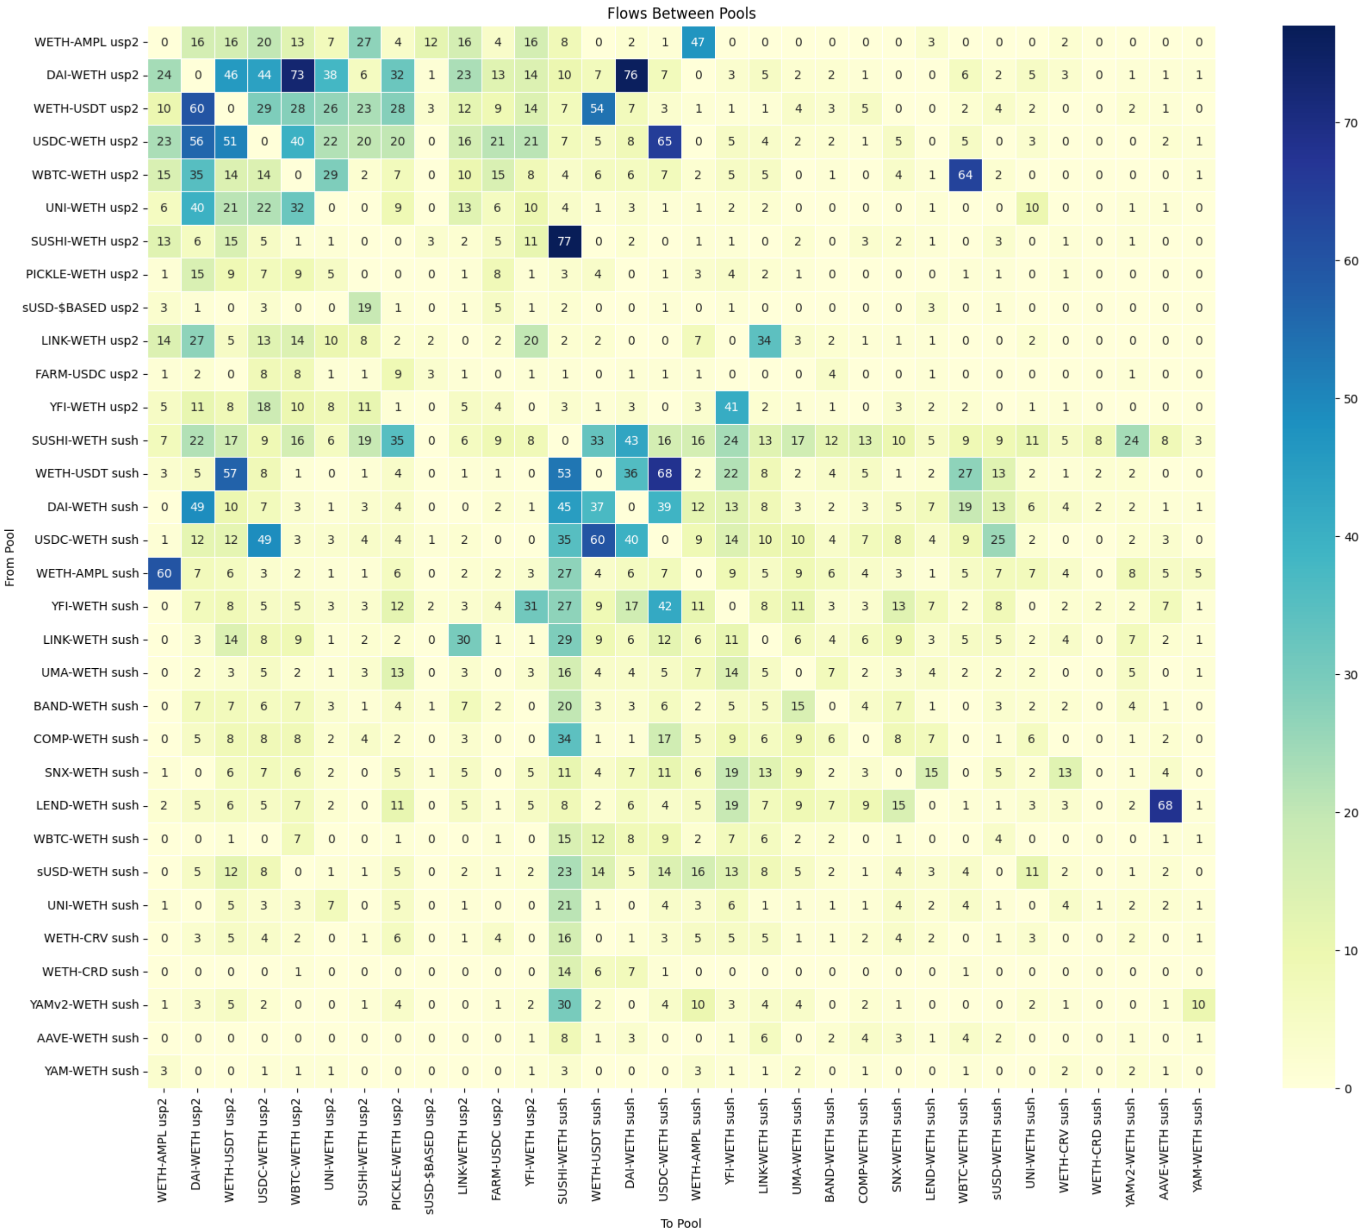

Fig. 18. Expanded liquidity graph topology given below where the number of flows is represented by the color gradient where darker tiles are the most connected. Data is collected between the top 20 liquidity pools on Uniswap and Sushiswap from September to December of 2020. The corresponding WETH/USDC pools exhibit overlap, but notably many other pairings as well. Sushiswap's exchanges reward token SUSHI is especially connected across multiple pairings, which is related to their vampire attack on Uniswap.