



Delft University of Technology

Decentralizing components of electronic markets to prevent gatekeeping and manipulation

de Vos, Martijn; Ishmaev, Georgy; Pouwelse, Johan

DOI

[10.1016/j.elerap.2022.101220](https://doi.org/10.1016/j.elerap.2022.101220)

Publication date

2022

Document Version

Final published version

Published in

Electronic Commerce Research and Applications

Citation (APA)

de Vos, M., Ishmaev, G., & Pouwelse, J. (2022). Decentralizing components of electronic markets to prevent gatekeeping and manipulation. *Electronic Commerce Research and Applications*, 56, Article 101220. <https://doi.org/10.1016/j.elerap.2022.101220>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Decentralizing components of electronic markets to prevent gatekeeping and manipulation

Martijn de Vos^{*}, Georgy Ishmaev, Johan Pouwelse

Delft University of Technology, Van Mourik Broekmanweg 6, Delft, 2628XE, Zuid Holland, The Netherlands

ARTICLE INFO

Keywords:

Decentralization
Gatekeeping
Market manipulation
Distributed ledger technology
Blockchain
Decentralized finance

ABSTRACT

The landscape of electronic marketplaces has been monopolized by a handful of market operators that have accumulated tremendous power during the last decades. This trend raises concerns about fairness and market manipulation by these operators acting as gatekeepers. These concerns have recently been outlined in the EU Digital Markets Act (DMA).

In this work, we highlight how technological logic of separation understood in the framework of decentralization can address manipulation concerns. As a first step, we devise a reference model of electronic marketplaces, containing six functional components, and outline how control over these components enables different manipulative practices by gatekeepers. We identify two dimensions of decentralization that can counterbalance monopolistic abuse of marketplace components. We then present a software implementation of our reference model and demonstrate how decentralization and unbundling of market components can alleviate manipulation and fairness concerns. We end our work with a review of related approaches and conclude that modular and interoperable marketplaces can enable an open ecosystem of fair electronic markets envisioned by the DMA.

1. Introduction

The significance of electronic markets and their infrastructures is hard to underestimate, given the spectacular growth of e-commerce during the last two decades. This growth, however, is a multifaceted trend that raises novel challenges and issues. One prominent concern is the monopolization of electronic marketplaces by “Big Tech” companies, such as Amazon, Google, Facebook, and Uber (Barwise and Watkins, 2018; Jullien and Sand-Zantman, 2020). At the one hand, concentrated control of electronic marketplaces can provide significant benefits for participants, e.g., reduced friction, easy market entry, and access to a large pool of potential buyers and sellers. On the other hand, several factors also invite closer critical scrutiny of this model for electronic marketplaces. Monopolization is not exclusively an issue of anti-competitive behaviour, a problem traditionally associated with monopolies. Instead, it is a much broader trend where a combination of network effects, economies of scale, and big data collection has made “winner takes all” a dominant strategy for companies operating electronic commerce platforms. This strategy prioritizes market domination through vertical and horizontal integration, consumer lock-in, and aggressive stifling of competition (Khan, 2018). The privileged

position of monopolistic market operators raises concerns around the fair treatment of participants, fair pricing, data privacy, corporate control of critical digital infrastructures, taxation, labour regulations, and manipulation of consumers and markets (Barwise and Watkins, 2018; Schechner, 2021; Göldi, 2020).

These concerns were recently laid out in the “Digital Markets Act” (DMA) regulation, proposed by the European Commission, that aims to address the practices of *gatekeeper* platforms who abuse their privileged intermediary positions. These unfair practices warrant a closer attention to the ways these platforms bundle together their core services and underlying infrastructures, and motivates the need for approaches to separate these components (EC, 2020b). *Structure separation*, the forcible separation of vertically integrated companies, has its limitations since application and monitoring of effective separation in quickly evolving high-tech markets can be undermined by obsoleting regulatory tools (Khan, 2019).¹ The EU expert report on DMA provides an opinion that not only legal but also *technological separation* of platform services from the infrastructure may be necessary to address fairness issues in electronic marketplaces (Cabral et al., 2021) (p.30). Currently, there is no comprehensive technical analysis on the feasibility

^{*} Corresponding author.

E-mail address: m.a.devos-1@tudelft.nl (M. de Vos).

¹ Historically, separation regimes were applied in particular markets and services where a bottleneck facility served as infrastructure or a critical intermediary in railroad, banking and telecommunication sectors.

of technical separation. We provide such an analysis and argue that technological separation, which we define as a decentralized design of electronic markets, brings the capability to address core issues of fairness and market manipulations by market gatekeepers.

An increasing amount of research provides new insights into practical solutions to disintermediate electronic marketplaces and replace centralized components with decentralized solutions such as distributed ledgers (Subramanian, 2017a). Thanks to significant advancements of blockchain technology, decentralized marketplaces is arguably one of the most prominent examples of this approach, illustrating a direction of design for electronic markets (Subramanian, 2017b).² The explosive growth of “Decentralized Finance” (DeFi) during the past two years is another vivid illustration (Chen and Bellavitis, 2020). DeFi is an experimental form of finance where financial products and assets are managed, traded, and lent using blockchain technology, thus avoiding trusted intermediaries or centralized coordination.

The idea of decentralized marketplaces is hardly a radical or novel proposal though. At the dawn of e-commerce, it has been pointed out that different designs of electronic markets could be exploited by the providers of technical infrastructures to capture customers in a system biased towards a particular supplier (Malone et al., 1987). At the same time, it has been argued that electronic marketplaces can facilitate unbiased markets enabled by electronic brokerage. Specifically, electronic matchmaking solutions were envisioned to create open, fair, and competitive markets on the Internet (Malone et al., 1987; Trastour et al., 2003). We agree with this analysis and argue that the decentralization of key marketplace components can address the root causes of market manipulations by intermediaries: *vertical and horizontal monopolization* of electronic marketplaces. We also suggest that this approach can take us one step closer towards an ecosystem of interoperable electronic markets — an ambitious goal of the DMA.

Paper outline and contributions

As a first step, we devise a reference model for separation logic in Section 2. Our model is based on the functional decomposition of an electronic marketplace and comprises six key components. It identifies necessary functional components of a generic marketplace corresponding to the business phases of market transactions. Our reference model adds new level of details to the understanding of gatekeepers marketplace platforms, and can be helpful both for researchers of decentralized marketplaces and policymakers dealing with the issues of gatekeeping in marketplaces. Our reference model also provides a vision of the possible key components for interoperable marketplaces where different service providers (e.g., matching or payment services) compete with each other.

In Section 3 we look into known examples of marketplace manipulations by market operators controlling key functional components of electronic marketplaces. We consider manipulative practices as defined in DMA: the possibility for the intermediary to exploit one side of a market and subsidize another. We have conducted an analysis of relevant research literature and produced a taxonomy of different types of manipulation over market participants. Using our reference model, we map each manipulation practice to one or more components that enabling these manipulative practices.

We find that vertical and horizontal monopolization of marketplaces, and abuse of “gatekeepers’ power”, lies at the root of these

² It is important to note here that the concept of a decentralized market can be employed in a different sense. As an economic-theoretical concept, a decentralized market refers to an abstraction of a market where participants randomly engage in bilateral contracts in the absence of a centralized dealer(s). This paper uses the concept of a decentralized market in a practical sense as a technological infrastructure that facilitates disintermediated trades between market participants.

manipulations. To address both concerns, we consider two different dimensions of decentralization logic: (1) on the high level of marketplace architecture, and (2) on the level of each of the identified marketplace components. To illustrate the viability of our reference model, we present a software implementation in Section 4. Our proof-of-concept implementation, named *AnyDex*, includes interactions between the functional components of a decentralized electronic marketplace.

In Section 5 we demonstrate how the decentralized architecture of AnyDex provides guarantees against the identified manipulation concerns listed in Section 3. We also show the resistance of our decentralized marketplace against more generic forms of manipulations, namely market information manipulation and counterparty fraud.

In Section 6 we provide a detailed review of current state-of-art solutions for the components of decentralized marketplaces, and argue that recent advancements in De-Fi technologies strongly suggest viability of this approach for marketplaces beyond financial products. We evaluate affordances and limitations of existing approaches, with a particular focus on scalability and interoperability.

We conclude our work with an analysis of current obstacles that need to be addressed in order to enable interoperable and open ecosystems for fair electronic marketplaces.

2. Key components of electronic marketplaces

To unbundle and decentralize key enabling components of electronic marketplaces, we first have to choose an appropriate framework for the identification of components comprising an electronic market. This is non-trivial as electronic marketplaces can be decomposed from different perspectives, e.g., from a business, legal, or technical perspective. Additionally, electronic marketplaces can be devised to enable different application-specific markets and therefore can include context-specific components. The DMA hints at a possible direction, highlighting some key enabling elements for all platforms. Namely, in parts that ban platform operators from imposing their own “identification services” on users (Art. 5e), and a ban on the bundling of core services together as necessary condition to the platform access (Art. 5f) (Comission, 2020).

It is helpful to consider that vertical integration is a core aspect of the platform monopolization problem. Many monopolistic platform operators exhibit properties of vertical integration either controlling several stages of the value chain, or by strategically positioning themselves between the tiers of supply chain Alt and Zimmermann (2019).³ A marketplace operator controlling key functional components for the different phases of marketplace trade essentially controls several stages of the value chain, ranging from the initial on-boarding of market participants to the finalization of orders. Generally, these phases are managed by different functional elements of the marketplace. We aim to identify key enabling components common to all electronic marketplaces, regardless of the goods and services being traded, or mechanisms used.

From a technological point of view, it is not trivial to disentangle these points of control given that these are opaque proprietary platforms designed and implemented as bundled services. To overcome this obstacle we consider separate functional components of a marketplace architecture that can be operated independently of each other. Specifically, we envision that each identified component can be realized with a software solution and as such, they can be operated by different parties ensuring that no single operator can concentrate control over all key enabling components. Our model considers technological separation of marketplace infrastructure as architecture-level decentralization.

³ The concept of vertical integration is usually characterized by the situation where a company (operator) takes complete control over several stages in the production and distribution of the product or service “value chain”.




Fig. 1. Our reference model for electronic markets, comprising six key components.

As a first step, we consider two most widely referenced approaches to the identification of components for electronic marketplaces. Interestingly enough, earlier conceptual models of electronic marketplaces have not experienced much transformation in the past decades.⁴ The reference model proposed by Schmid et al. is still very much a standard categorization for the elements of contemporary electronic markets (Schmid and Lindemann, 1998a). Their reference model includes business models, business scenarios and technological elements. They identify three business phases: (1) the *information phase* during which suppliers and consumers acquire information about potential market partners as well as goods and services; (2) the *agreement phase* during which the trade conditions are negotiated, potentially resulting in a contract; and (3) the *settlement phase* during which the involved parties fulfil the agreed-upon terms of the contract. In this paper, however, we are primarily interested in the technological elements, which can be categorized as enabling components corresponding to specific business phases of market transactions.

Alternatively, Reich et al. propose an architecture for electronic stock markets (Reich and Ben-Shaul, 1998). They take a system point of view and abstract a high-level generic architecture. Their decomposition includes three main parts: (1) the *front-end* that defines the rules for the acceptance of incoming orders, (2) the *trading floor* that is tasked with processing incoming and includes components for timing, matching and prioritizing orders, and (3) the *back-end* that informs traders when the status of their order changes. We agree that high-level system abstraction is the appropriate approach to grasp key enabling components of a marketplace. However, the analysis by Reich et al. provides only a basic schematic view of the system and does not consider interactions between the components (Reich and Ben-Shaul, 1998).

In Fig. 1 we visualize our reference model for electronic marketplaces, comprising six functional components. These components are the result of our literature analysis on both traditional electronic marketplaces and recent innovations in decentralized marketplaces, e.g., blockchain-based solutions. In the remainder of this section, we elaborate on each component and describe how they are realized in traditional electronic markets.

2.1. Information management

Electronic markets require a mechanism to manage and store all information associated with the market. Early designs of electronic marketplaces mainly focus on this specific component, e.g., the Electronic Product Catalogue (Schmid and Lindemann, 1998a). However, market information not only includes product listings and pricing information but also includes outstanding orders, profile information of participating traders, and details on their historical transactions. The latter can, for example, be used to estimate the trustworthiness of

market participants before engaging in a trade. In Fig. 1, information management provides a communication basis for the other five components.

Except for a few notable proposals in academic research (e.g., GEM (Rachlevsky-Reich et al., 1999a)), centralized information management solutions were widely considered to be the only viable approach until the recent emergence of blockchain-enabled marketplaces. This is a dominant strategy to date: all major electronic marketplace platforms take a centralized approach to information management. In such marketplaces, information is stored and managed on servers under the control of a single market authority.⁵ A centralized approach to the management of market information has three key benefits. Firstly, it can enable high-speed access to relevant information by traders since the market operator can optimize their infrastructure, e.g., by installing fast uplinks or applying geo-distribution techniques to reduce latency. Secondly, since there is a centralized point of control, it is easier to address targeted attacks and filter out invalid information, e.g., when a trader spams the market with invalid orders. Thirdly, centralized servers are relatively straightforward to set up, operate, and maintain by the market operator from a technological point of view. We also identify two disadvantages of centralized information management. Firstly, information management systems of market operators and associated protocols are often proprietary software and not open for inspection by traders. Secondly, they provide a single point of failure where even a minor configuration error can result in prolonged market downtime (Schneider et al., 2021).

2.2. Identity management

On all electronic marketplaces a participant interacts with the market using at least some form of a digital identity. We outline five purposes of an identity layer in the context of markets and online economic activity. First, identity management is required to onboard market participants and to associate orders, transactions, and trade activity with a participant. Second, it ensures accountability of one's actions within the market in case of a dispute between a buyer and seller. Third, it prevents the situation where a user can easily re-enter the market under a different identity after having committed fraud. Forth, identity verification is often a part of the regulatory compliance of market operators, as often required by anti-money laundering policies imposed by governmental bodies.⁶ Fifth, identity management is a necessary enabling component for other marketplace components, such as reputation mechanisms, risk management, and conflict resolution.

⁵ We note that market operators using centralized information management can internally manage market information in a *distributed* manner where the storage functionality is dispersed over multiple servers, e.g., geographically distributed, but not decentralized in terms of access and control.

⁶ It is interesting to note, however, that contraband “dark” marketplaces operating outside any jurisdictions are also functionally dependent on identity management mechanisms, such as pseudonymous identities and reputation (Tzanetakis et al., 2016). This suggests that identity mechanism is a functional component.

⁴ This observation indirectly supports the hypothesis that the growth of proprietary monopolistic e-commerce platforms in the past twenty years has not contributed to the available academic research in this field (Azevedo and Weyl, 2016). Detailed analysis of this argument, however, lies outside the scope of our work.

2.3. Risk management

Risk management is a key component of any market. Participants should take into consideration *counterparty risk*, which is the risk associated with one of the trading parties defaulting on contractual obligations. Management of counterparty risks is crucial since in the absence of risk mitigation mechanisms, market participants are reluctant to engage in trades. In a functional sense, risk management may involve intermediaries for settlement, contingent contracts, reputation, collateral deposits, or market specific insurance pool which can compensate losses for trader exposed to counterparty risks. Risk management is a process that happens before settlement.

In finance markets, risk management is usually carried out by central clearing parties, for example, by a company that estimates credit risks of market participants. This solution, however, may introduce its own systemic risks (Menkveld, 2017). It needs to be noted that even though in practice risk management is mostly associated with financial markets, at the high level of market abstraction any type of trade involves counterparty risks. Generally speaking, counterparty risks are necessitated by information asymmetries, since in the presence of perfect information, market participants can choose between counteragents and rationally reduce the size of risk exposure, obviating the need for risk management mechanisms (Stephens and Thompson, 2017). We argue, thus, that not only financial but many other types of emerging and novel markets such as sharing economy markets, introduce new dimensions of information asymmetries and thus require dedicated mechanisms for risk management. For example, Airbnb enables prospective guests to estimate the trustworthiness and reliability of a host through several means, e.g., photos, reviews, and performance indicators like response time.

2.4. Matchmaking

Matchmaking between buyers and sellers is a prerequisite for online trade and is essential for any two-sided marketplace (Veit et al., 2002).⁷ Matchmaking is defined as the process of mediating supply and demand in markets, based on profile information (Veit et al., 2002). This process depends on the individual constraints and preferences of market participants.

In financial markets, matchmaking is often an automated process where an algorithm matches incoming buy and sell orders according to a matching policy. A common matching policy is price-time, where orders are first matched based on their price and then based on their creation time (Mavroudis and Melton, 2019). In other markets, however, matchmaking is a manual process where users get access to the full catalogue of orders, e.g., products. This approach is common for marketplaces acting in the sharing economy, such as Uber and Airbnb, where users can choose their preferred counterparty. Manual matchmaking can be streamlined with a recommendation, searching, and filtering engine.

In earlier reference models, matchmaking was delegated either to market makers (Reich and Ben-Shaul, 1998), or dedicated intermediaries, acting as mediating electronic product catalogues (Schmid and Lindemann, 1998a). However, recent advancement in matchmaking algorithms has transformed matchmaking into a key component of electronic marketplaces, fuelling a surge in “matching markets” (Azevedo and Weyl, 2016). We suggest thus, that matchmaking mechanisms in electronic markets should be properly analysed as a stand-alone functional element. Matching algorithms at proprietary platforms are often bundled together with other core elements.

Based on our empirical research and engineering experiments, we can distinguish three different types of matching solutions: centralized, federated and decentralized matchmaking (see Fig. 2). Centralized

matchmaking (Fig. 2(a)) is the most predominant approach where the market operator maintains centralized infrastructure to match incoming orders. With federated matchmaking (Fig. 2(b)), other parties can act as matchmakers and trader send their orders to a single matchmaker. The main idea behind decentralized matchmaking (Fig. 2(c)) is that a single order can be sent to multiple matchmakers simultaneously, and can be shared amongst them.

2.5. Settlement

Settlement is the process of fulfilling the agreed-upon obligations by trading parties (Schmid and Lindemann, 1998b). Settlement usually proceeds after the matchmaking process when two parties have created and signed a contractual agreement. Before settlement, prospective traders can negotiate agreements about the conditions of the upcoming date, e.g., the delivery date and pricing of goods. It needs to be noted that settlement is highly context-specific and may vary greatly across different types of markets. Compared to consumer markets where settlement involves deliveries, settlement on financial markets can involve the processing of trades by specialized entities such as a clearinghouse. Thus, it is common practice in electronic marketplaces to have a trusted intermediary carrying out the settlement process (Giaglis et al., 2002). This trusted intermediary can, for example, take care of payment processing, crediting, or the transportation of goods.

2.6. Conflict management

As the scale and reach of electronic markets expands, it becomes more likely that conflicts and disputes arise between participants regarding execution of orders and performance of trades. From the functional point of view, a mechanism addressing these issues can be understood as a mechanism of technology-augmented dispute resolution.⁸ Dispute resolution can occur either during or after settlement.

Earlier models of electronic markets outsourced conflict management to external parties or intermediaries (Schmid and Lindemann, 1998b; Reich and Ben-Shaul, 1998). However, various major market operators have been developing in-built conflict management solutions. eBay’s “Money Back Guarantee” is a type of integrated dispute resolution, acting as an insurance-type safeguard in case the buyer does not receive the ordered item or when the item does not match the listing description. This solution might also involve third parties as eBay may request the interfaced payment operator (e.g., PayPal) to withdraw the funds on the seller’s account to enforce this responsibility. eBay also experimented with creative process design, producing the “eBay community court” pilot in late 2008. Under the community court scheme, sellers disputing negative buyer feedback could submit their complaint to a randomly selected panel of jurors. This approach represented one of the first attempts to use crowdsourcing in dispute resolution.

3. Manipulation through the control of market components

In recent years we have seen a transformation of marketplaces towards integrated platforms controlled by a single intermediary. Such marketplaces not only provide front-end services in a form of a website or a mobile application, but also a back-end with the necessary components described in Section 2, provided by the same intermediary. This uniquely privileged position of platform owners, who design, implement, and operate these components as opaque proprietary solutions

⁸ Online Dispute Resolution (ODR) sometimes can also refer to technological tools used to assist in legal proceedings and litigations, here we use the term of dispute resolution only to refer to governance and technological mechanisms used to resolve market conflicts occurring in regards to the fulfilment of contractual obligations between market participants.

⁷ Matchmaking is sometimes referred to as *brokering*.

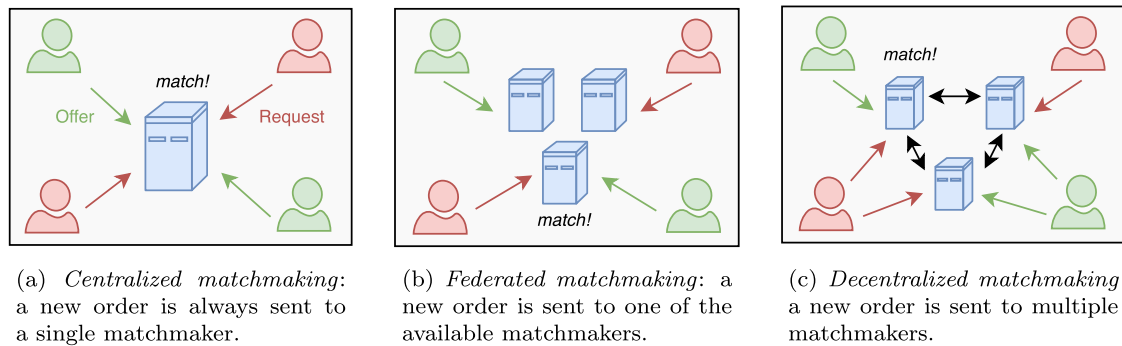


Fig. 2. Three approaches for matchmaking. Traders create offers and requests (coloured green and red respectively), which are matched by matchmakers (depicted in blue). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1

Manipulation practices by the operator in electronic markets and the required control of components (abbreviated) to perform the manipulation (● indicates that the control over the single component is enough for this type of practice, ● indicates that control over this component is needed together with other components, and ○ indicates that it is helpful to have control over this component but not necessary).

Manipulation	Component					
	Information Management	Identity Management	Risk Management	Matchmaking	Settlement	Conflict Management
Front running	●			○		
Delaying or prioritizing orders				●		
Price steering	○			●	●	
Price discrimination	●					
Trader impersonation		●				
Censorship	●	○		○		
Behavioural manipulation	●					
Quality filtering	○		●			●

enables various kinds of manipulations. For example, market operators can abuse information systems, software implementations, network communication, and algorithms for selfish interests.⁹

It is safe to say that no single definition of manipulation can grasp the variety of issues associated with unfair marketplace practices. On the one hand, there is a gap between legal definitions of fairness in competition law that emphasize principles of non-discrimination of market participants. On the other hand, there is a somewhat broader issue of an unequal distribution of costs and benefits in economic sense. This broadness, however, does not preclude us from operationalizing some narrower definition specific to the context of marketplace platforms. We follow the definition of unfair platform operator practice proposed by the EU expert report on DMA (Cabral et al., 2021). It applies in a case where an intermediary or gatekeeper mediates a three-party trilateral exchange, with a possibility to *exploit one side of a market and subsidize another*. Given the opaque nature of centralized electronic commerce platforms built on proprietary systems, we often have only inferential or circumstantial evidence of gatekeeping abuse and manipulative practices. There is, however, a growing understanding of these

practices, being revealed through reverse engineering (Agmon Ben-Yehuda et al., 2013), or lawsuits (EC, 2020a). Thus, we have conducted an analysis of relevant research literature and produced a taxonomy of different types of manipulation over market participants. We do not suggest that our analysis necessarily implicates market operators in the manipulative practices. Rather, our goal is to identify enabling factors, which is consistent with *per se* approach of DMA. More specifically, we aim to understand how the control over one or several functional components enables manipulation.¹⁰

Table 1 lists different forms of manipulations by the market operator. We show over which component the market operator requires control to succeed in a particular manipulation effort. These manipulations are not mutually exclusive, and one form of manipulation can further increase the success of other manipulation efforts. This analysis also shows that the control of several key market components highlighted in Section 2 enables a wider scope of manipulative practices than just control of a single component. Our analysis reveals that information management is a critical component and control over it enables a variety of manipulations such as front running and censorship. This is not an exhaustive analysis but it demonstrates that manipulation is deeply intertwined with the vertical integration of a marketplace in the hands of a gatekeeper.

3.1. Front running

Front running is the situation where a participant acts on inside information to get a time advantage when responding to pending orders of other traders. Particularly, it is an issue in financial markets where orders are automatically matched and executed. In a stock exchange, for example, a broker could front run on one of their clients' order. This is, however, considered an illegal practice in many markets (Lin, 2016). We refer the reader to the work of Markham et al. for a regulatory and historical perspective on front running in electronic markets (Markham, 1988). Front running is enabled when an operator has priority access to incoming market information. Additionally, a market operator can increase the chances of successfully front running on a particular order by leveraging its control over matchmaking by delaying orders that would undermine the front running attempt.¹¹

While front running is typically associated with financial markets, the development of algorithmic pricing mechanisms and recommender algorithms have made front running also possible in other types of markets. For instance, the recent anti-trust investigation of Amazon by the European Commission suggests that a market operator controlling

⁹ This is not to suggest that intermediaries are the only culprits of unfair markets, but in the scope of this research we specifically focus on service providers and omit manipulation concerns stemming from traders themselves, e.g., exploiting technical characteristics of market algorithms. For some identified manipulations, however, traders can collude with the market operator to get an unfair advantage over other traders (Mavroudis, 2019).

¹⁰ The DMA's prescriptions are *per se* rules that "apply independently of actual, likely, or presumed effects of the conduct of a given gatekeeper".

¹¹ Front running also occurs in blockchain-based marketplaces where a miner acts on incoming transactions before they are stored on the distributed ledger.

information management can conduct front running for any type of traded goods. By dynamically monitoring relevant data on user behaviour, a market operator has the ability to recommend a potential buyer goods at lower prices before checkout, thus front running other sellers on the platform. It is suspected that Amazon exploits the “Buy Box” by frequently offering its own products to customers (Scott and van Dorpe, 2020).¹² This form of front running can be problematic if the platform operator also acts as a seller within the same market system, but it can also manifest as a result of collusion between a market operator and certain sellers.

3.2. Delaying or prioritizing orders

Market operators can delay or prioritize particular buy or sell orders when they control the matchmaking logic. A market operator can, for example, prioritize its own orders to increase economic gains (e.g., by front running orders) or defer the execution of orders to manipulate the market price of a particular asset and to attract trading volume. Since many financial markets have to process orders within milliseconds, and because most order matchmaking implementations are proprietary, this manipulation is challenging to detect. Order manipulation is not a problem exclusive to financial markets: it is suspected that the ride-hailing platform Uber actively manipulates the matchmaking process between passengers and drivers where passenger satisfaction is preferred over the satisfaction of drivers (Bokányi and Hannák, 2020).

Though we focus on manipulation efforts carried out by market operators, traders can also exploit the technical characteristics of match-making engines. We refer the reader to the work of Mavroudis et al. for a thorough discussion on this topic (Mavroudis, 2019).

3.3. Information manipulation

Market operators can manipulate the interaction between sellers and buyers by abusing control over the flow of market information. We discuss two forms of manipulations by the market operator that affect the price of products being traded on the market: price steering and price discrimination.

3.3.1. Price steering

Among many examples, Lyft controls pricing to facilitate rides between riders (buyers) and drivers (sellers), Airbnb controls search results and recommends pricing to influence matching between hosts (sellers) who rent their houses to guests (buyers), and LendingClub assigns a creditworthiness score to borrowers (sellers) who are applying for a loan from investors (buyers) (Pavlov and Berman, 2019). Because sellers set their prices based on their beliefs about buyer demand, the platform can influence competition and price levels by supplying artificial information to sellers or by modifying prices directly. A market operator can also steer the price of assets if it controls the settlement process and intervals (Jun and Rui, 2013).

Because of its role as an intermediary, Uber, for example, has access to significant data unavailable to both drivers and riders and a capacity to monitor, which is not reciprocally available to either group. By utilizing this information asymmetry, the platform can leverage “access to information about users and their control over the user experience to mislead, coerce, or otherwise disadvantage sharing economy participants” — a claim reflected in the revelation of Uber’s manipulation of drivers (Bamberger and Lobel, 2017).

3.3.2. Price discrimination

Price discrimination occurs when two market participants are shown inconsistent prices for the same product. For example, it has been argued that Uber’s reliance on discrete surge areas introduces price discrimination into their system: two users standing a few metres apart may unknowingly receive dramatically different surge multipliers. For example, 20% of the time in Times Square, customers can save 50% or more by being in an adjacent surge area (Chen et al., 2015). Chen et al. argue that Uber’s reliance on black-box algorithms makes their system more vulnerable to manipulation than other online marketplaces. Uber operates as a black-box: they do not provide accurate data about supply and demand, and an opaque algorithm sets surge multipliers. This lack of transparency has led to concerns that Uber may artificially manipulate surge prices to increase profits, as well as apprehension about the fairness of surge pricing. Dynamic pricing algorithms can implement collusive strategies that harm consumers. The US Justice Department successfully prosecuted several individuals who implemented a price-fixing scheme on Amazon using algorithms (Chen et al., 2016).

Given the structure of the platform as a multi-sided peer network, which means that both supply and demand are highly elastic, dynamic pricing can enable platforms to engage in practices approaching first-order or “perfect” price discrimination, by which a firm personalizes prices to reflect the maximum an individual is willing to pay (Bamberger and Lobel, 2017). Uber’s surge pricing model has been criticized as exploitative of consumer’s willingness to pay more in times of bad weather or increased demand. Algorithmic pricing, used by a platform with market dominance and the capacity to amass significant personal data about individual consumers, could engage in behaviour that could approach perfect price discrimination: person-specific pricing, that charges each user their “exact reservation price” — the maximum they would pay.

3.4. Censorship

Censorship is a multi-faceted issue in e-commerce platforms covering multiple aspects of the market process. Not all types of censorship necessarily constitute market manipulation or are morally problematic. For instance, we could consider relatively uncontroversial instances when it can be desirable to prevent markets that may be considered ethically unacceptable or repugnant (e.g., human trafficking or trading weapons of mass destruction). However, outside such contexts market censorship may take forms of manipulative practices whenever omission of information or goods is not only legally and ethically arbitrary but also brings economic benefits for the market operator.

One example of such manipulative censorship is the possibility of collusion of a market operator with other market participants. Such a market manipulation is illustrated by the removal of *GameStop Corp.* stocks from the list of tradable items by the operator of the Robinhood app.¹³ The market operator can leverage its control over information management to hide information of particular traders (Duggan, 2021). This is not the only type of manipulative censorship by market operator. When a market operator controls identity management, it can engage in arbitrary censorship, preventing particular traders from participating in the market to reduce competition. Additionally, when a market operator control the matchmaking process, it can ignore the orders of particular traders.

¹² Amazon is estimated to offer 75% of the best selling goods on its platform (Scott and van Dorpe, 2020).

¹³ While there is no evidence that this removal was a direct result of collusion between the Robinhood operator and other market participants, this decision did result in a significant price drop of “GameStop Corp.” stock which economically benefits large traders indirectly affiliated with Robinhood.

3.5. Behavioural manipulation

Behaviour manipulation by platform operators has recently gained much attention in consumer markets and by researchers on “dark patterns” in user interfaces (Narayanan et al., 2020). Conducting this manipulation requires control over information management. It may be argued that attempts to influence the market participants’ decision-making process do not necessarily constitute malicious manipulation. In fact, most marketing and advertising techniques as old as markets themselves use behavioural manipulation. However, in the context of electronic markets characterized by immense information asymmetries between market participants and operators, these practices take qualitatively new forms. Firstly, control over the participants’ primary interaction paths in online services or websites allows platform operators to steer and nudge market participants. Secondly, control over marketplace information flows allows platform operators to evaluate and improve the efficacy of these methods at scale. While large scale empirical studies on these techniques are still limited, researchers have discovered dark patterns interfaces at around 11.1% websites in the dataset of 11K shopping websites (Mathur et al., 2019). The effects of these manipulative practices can vary from nudging market participants into suboptimal market transactions to the coercive imposition of unwanted contracts (e.g., hidden subscription services or adding products to a shopping cart without user’s consent).

3.6. Quality filtering

Review manipulation is a key concern on electronic marketplaces (Mayzlin et al., 2014). This problem, however, is only a specific facet of a more general issue of market design. Feedback and review systems streamline risk management and conflict resolution procedures (Bolton et al., 2018). On the one hand, reviews can help market participants to address counterparty risks through the reduction of information asymmetries. On the other hand, strategic submission and withdrawal of reviews can be used as a tool for post-settlement dispute resolution between market participants. Thus, such mechanisms must satisfy neutrality and impartiality criteria. Implementation of these mechanisms by platform operators can create a conflict of interest between their economic interests and quality filtering of goods and services provided at the platform. While intuitively, it may seem that market operators are always interested in facilitating high-quality trades, economic modelling suggests that to maximize their intermediary fees, operators of peer-to-peer markets should always aim to increase the number of sellers on the platform (Pavlov and Berman, 2019). This is consistent with empirical findings, suggesting that while Amazon takes significant efforts to remove fake reviews from its website, sellers using fake reviews and offering low-quality products are never penalized in any way (He et al., 2020). This phenomenon suggests that platform operators controlling risk management and conflict management mechanisms can take a strategic approach toward maximizing their intermediary fees at the expense of quality filtering.

4. AnyDex: A Decentralized and Manipulation-Resistant Marketplace

In Section 2 we have devised a reference model that considers separation as decentralization at a system’s architecture level. We then identified in Section 3 various manipulation practices of markets by platform operators. Separation of functional components, however, is not sufficient to address all identified types of manipulation. From Table 1 we observe that for some types of manipulation it is sometimes sufficient to control a single component, e.g., information management. If in a certain market one of few market operators are dominant, this can still enable various manipulative practices. We suggest considering two dimensions of decentralization for an electronic marketplace, along the lines of *horizontal and vertical monopolization*.

Vertical monopolization can be addressed at the level of architectural decentralization, where no single operator has control over multiple functional components. Counteracting horizontal monopolization requires decentralization at the level of individual components.

To illustrate the feasibility of market decentralization at both of these levels, we present *AnyDex*, our decentralized and manipulation-resistant marketplace that enables generic trade at scale. AnyDex is based on the principle of decentralizing all components of electronic marketplaces (see Section 2) and also addresses various manipulation concerns associated with centralized control over each of these components (see Section 3). AnyDex combines five decentralized mechanisms designed by the Delft Blockchain Lab, namely the TCID identity system (Stokkink et al., 2021), the TrustChain distributed ledger (de Vos and Pouwelse, 2020), the MATCH matchmaking algorithm (de Vos et al., 2020), the NetFlow reputation mechanism (Otte et al., 2020), and the XChange trading protocol (de Vos et al., 2021). Each mechanism has been implemented, evaluated, and deployed. For a more extensive analysis of each mechanism, we refer the reader to the original papers describing the integrated systems. While the focus of this research is a decentralization of market components to prevent market operators’ manipulation, we acknowledge that manipulation by traders becomes much more important in decentralized environments. We address this feasibility issue in Section 5.9.

We have fully implemented AnyDex in the Python programming language and made its source code available in a GitHub repository.¹⁴ A key goal of AnyDex is to avoid centralized points of control and coordination while ensuring resistance against various forms of manipulation. As a result, all market components are decentralized and fully managed by traders themselves. AnyDex includes tools that enable system designers to leverage the integrated components. System designers can deploy their marketplace using AnyDex without permission of an authoritative party.

The AnyDex system architecture is visualized in Fig. 3. In the following subsections we describe the specifications of each component. Then, in Section 5, we outline how AnyDex addresses the manipulation concerns listed in Table 1.

4.1. Information management

AnyDex uses a peer-to-peer network for communication between traders. All market information, e.g., orders, trade details, and reviews, is stored on a tamper-evident, scalable, and distributed ledger. These two components make up the bottom layers in Fig. 3 (in grey) and are explained next.

4.1.1. Network layer

The network layer enables peer-to-peer communication between traders. We have implemented the network layer using an existing network library named IPv8.¹⁵

Each participant in AnyDex possesses a cryptographic keypair, consisting of a public and private key. The public key uniquely identifies the user in the network and the private key is used to digitally sign all outgoing messages. Since it is unlikely that each user knows the identity of all other users, a user connects to roughly twenty other users. More specifically, AnyDex maintains an *unstructured network* which resembles the structure of many peer-to-peer solution (Fletcher et al., 2004). When a new user joins the network, it contacts a bootstrapping service that notifies the newly joined participant of some random, existing users. This service is only used to bootstrap new users in the network.

¹⁴ See <https://github.com/tribler/anydex-core>.

¹⁵ See <https://github.com/tribler/py-ipv8>.

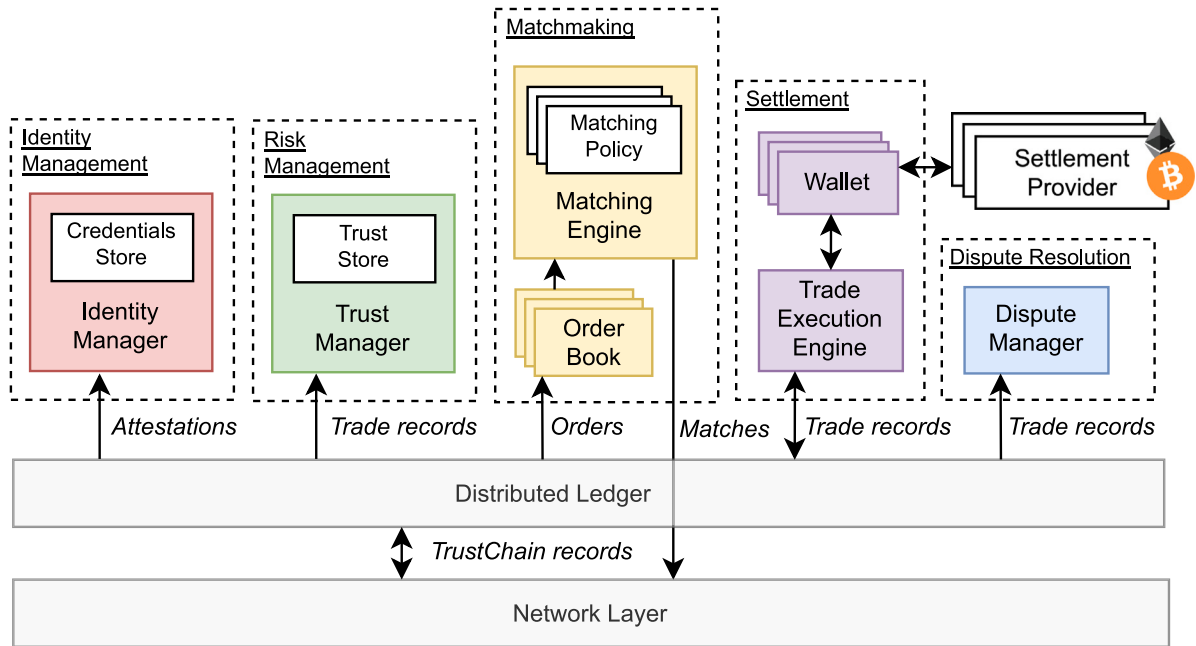


Fig. 3. The system architecture of AnyDex, a decentralized and manipulation-resistant marketplace.

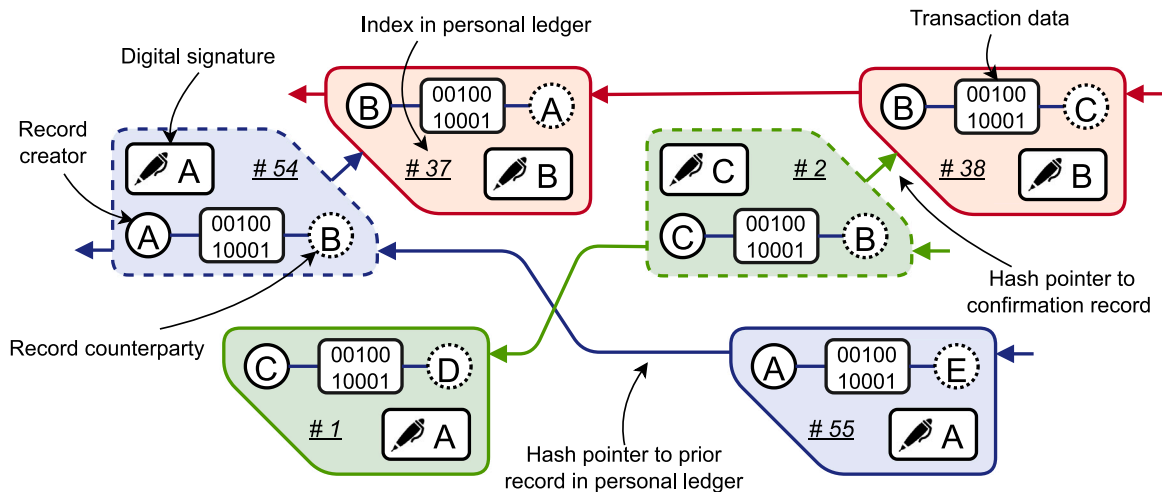


Fig. 4. The TrustChain data structure, used by AnyDex to store market information. Proposal and confirmation records have a solid and dotted border, respectively.

4.1.2. Tamper-evident distributed ledger

AnyDex stores all market information on a distributed ledger. For this purpose, we avoid using conventional blockchain ledgers, e.g., Ethereum or Bitcoin, where maximum achievable throughput can be a bottleneck. Most blockchain architectures require participants to perform a resource-intensive global consensus mechanism in order to avoid invalid transactions from being stored on the blockchain.

Instead, AnyDex uses a scalable, and lightweight ledger, named TrustChain, that is specifically designed for the tamper-evident accounting of information (de Vos and Pouwelse, 2020; Otte et al., 2020). Each user in AnyDex maintains a *personal ledger* with tamper-evident records. This is a key property of the TrustChain ledger and sets it apart from traditional blockchain architectures. These records contain market information and interaction data. Each record has a type field that signals what kind of information the record contains. By default, AnyDex supports the creation of records that indicate an order, trade, or review. These three operations are common across many market domains. At the same time, a system designer can deploy its market that contains custom record types. For example, a merchant might publish a

general notification about their stock with AnyDex. Except for the first record, each record in a personal ledger links to the previous one with a hash. A record can contain arbitrary data and can also point to records in the personal ledger of other users, resulting in a global graph, see Fig. 4. An interaction between two users is captured using a pair of a proposal and confirmation record. The confirmation record links to the corresponding proposal record. This structure enables users to quickly detect the modifications of market information after it has been stored, a process we describe in Section 5.

4.2. Identity management

To raise the barrier for fraud and identity-based manipulation by marketplace operators AnyDex integrates the modular and universal TCID identity management component. TCID is a self-sovereign identity (SSI) solution and a decentralized infrastructure for private data storage by identity holders, or *subjects*. Trusted third parties act as issuers and can attest particular claims created by subjects. Subjects can selectively disclose their claims and other parties can verify the validity of a claim.

A key aspect of TCID is data portability which allows traders to re-use their verified identity data in different AnyDex markets or on other platforms without the permission of the platform operator. TCID also provides a communication substrate for claims and verification data flows. Peers cannot impersonate each other in TCID system, as messages are signed by digital signatures from decentralized PKI (public key infrastructure), to ensure authenticity and integrity. Traders keep track of validated identities and refuse to interact with traders that have an unverified identity. This logic is implemented in the identity manager, included in Fig. 3. When a user a receives a message from another user b for the first time, a requests an attestation from b that forces b to prove its identity (without disclosing sensitive user data). Until b sends the requested attestation to a , a refuses to interact with b . As such, only participants that (1) have a signed attestation by a trusted third party, and (2) disclose this attestation to other traders, are able to interact and trade with each other.

Our prior TCID experiments measure the performance of the enrolment and verification of credentials, in terms of speed and network overhead (Stokkink et al., 2021). Different TCID implementations all finish within 3 s, which is on par or better than other existing verifiable credentials implementations. Furthermore, network traffic evaluation experiments suggest that even most demanding TCID protocol implementation utilizing zero-knowledge proofs requires only up to 70 kilobytes of data to be transferred per verification.

4.3. Risk management

AnyDex uses the NetFlow reputation mechanism to calculate a trustworthiness score of prospective trading counterparties (Otte et al., 2020). This helps traders to make an informed decision on prospective trading partners and to quantify the risks associated with interacting with unknown traders. The NetFlow mechanism uses collected TrustChain records created by traders. The AnyDex software builds and maintains a directed *trust graph* in which each vertex represents a user and each edge from u to v is weighted by weight w (where $w \geq 0$) by how much trust user u puts in user v . The weight on an edge (u, v) in the trust graph increases after a successful trade and decreases when party v commits fraud during the settlement phase of a trade (further discussed in Section 4.5). This trust graph is dynamically updated when new TrustChain records are received.

NetFlow is based on the notion of transient trust: if a user u trusts user v , and v trusts user w , then u also trusts w to a certain extent. More specifically, NetFlow uses a max-flow algorithm that computes the maximum flow between two users u and v . This flow gives an estimate on the trustworthiness of v according to u . Flow-based algorithms, although executed by a centralized party, have been explored before to estimate the risks in online marketplaces (Post et al., 2011). We note that if there is no path between u and v in the trust graph, NetFlow is unable to estimate the trustworthiness of the target trader. However, a trader can still use other means to estimate their trustworthiness, e.g., by reading the reviews left by other users. Trustworthiness scores are computed from the perspective of each trader and as such, different parties might assign differing trustworthiness scores to the same trader.

4.4. Matchmaking

AnyDex integrates the decentralized MATCH mechanism to match buy and sell orders (de Vos et al., 2020). The main idea of MATCH is that each trader in the AnyDex network acts as matchmaker for others by bundling storing incoming market orders in order books and by notifying others when a match for their orders has been found. New orders are sent to several available matchmakers. MATCH addresses fairness issues arising when matchmaking is performed by a single party, a common practice in many marketplaces. Specifically, MATCH makes it non-trivial to delay, prioritize, or hide orders (also see Section 5.2).

4.4.1. Order books

AnyDex bundles all buy and sell orders in an *order book*. The order book lists the specific assets that are being bought or sold within the market and provides traders with a view of the current supply and demand. AnyDex maintains distinct order books for different order types. Maintaining distinct order books at the same time enables order management with a single infrastructure across different trading domains.

4.4.2. Matching engine

The matching engine matches incoming offers (sell orders) and requests (buy orders) with existing orders in order books. In AnyDex, users themselves operate a matching engine, collect orders, and attempt to match new incoming orders, according to prescribed matching policy and protocol. The three steps of the matching process are: (1) Order Creation; (2) Order Negotiation; and (3) Settlement and Order Updating. These steps are visualized in Fig. 5. User digitally signs buy or sell order and sends it to a subset of available matchmakers in AnyDex. Valid offers and requests are matched against earlier received requests and offers, respectively. When a matchmaker has found two matching orders, it sends a *match notification* message to one of the order creators. The message is added to the user's match queue. Users then start a peer-to-peer negotiation process with a prospective counterparty, potentially leading to a contractual agreement and trade. When the trade is complete, the order creators inform their matchmakers. For a detailed analysis of the data flow we refer the reader to our prior work (de Vos et al., 2020).

4.4.3. Matching policies

The matching engine can contain multiple *matching policies*. A matching policy predicates whether an offer and request match, and is applied to a single type of order book. It takes a single offer and request as input and outputs whether these orders match, based on their specifications. The most common matching policy in financial markets is the *price-time matching policy*, where orders are first matched based on their price, and then based on order creation time in case of a tie-breaker (prioritizing older orders). AnyDex allows system designers to define custom matching policies for different order types.

4.5. Settlement

AnyDex integrates the XChange trading protocol for the settlement of trade (de Vos et al., 2021). XChange is specifically designed to trade assets between isolated ecosystems, e.g., different blockchains, without requiring a trusted intermediary that coordinates the trade. XChange accounts all trade activity on a distributed ledger, including contractual agreements, payments, and trade finalization. A trade is settled in three steps. First, prospective counterparties construct a *trade agreement* record, containing the specifications of the upcoming trade, and store it on TrustChain. Second, counterparties alternately issue *payments*, exchanging their assets or goods. Each issued payment is recorded by the initiator within a TrustChain record. The trade execution engine (see Fig. 3) keeps track of the current status of ongoing trades. Third, when all payments have been conducted, both traders confirm the *completion* of a trade on TrustChain. Traders manage their assets in AnyDex using integrated wallets (e.g., crypto-native assets) and use settlement providers to transfer these assets. Various types of trades may require services of different settlement providers, e.g., shipment companies, or banking services. In this case, settlement providers can confirm the completion of a payment in the form of a signed statement, e.g., a proof-of-delivery.

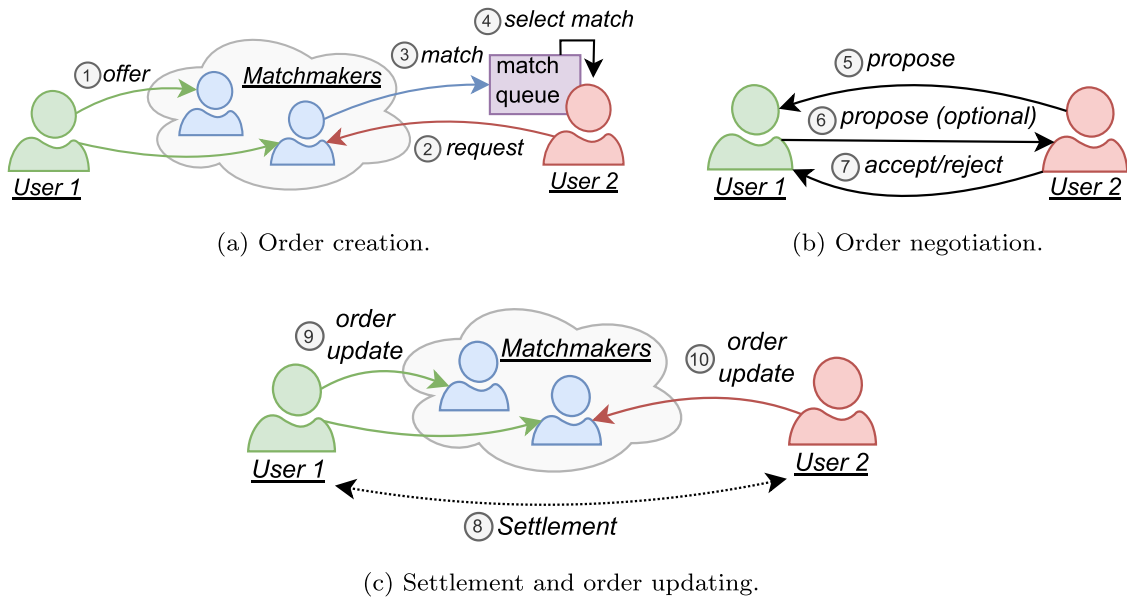


Fig. 5. The three steps of the MATCH matchmaking protocol.

4.6. Dispute resolution

By recording all trade agreements and payments, AnyDex allows any trader to detect disputes between traders. When a dispute arises between two traders, AnyDex provides the functionality for these traders to communicate and resolve the dispute without arbitration. When the conflict cannot be resolved amongst users themselves, an *arbitrator* in the AnyDex network is contacted that can investigate the dispute. Traders can indicate to act as arbitrator, and their availability and dispute verdicts are recorded on the TrustChain ledger. Arbitrators can coordinate with settlement providers, for example, to revert a payment. Different trading domains can have different arbitrators since they require an arbitrator to have domain knowledge.

5. Resistance of AnyDex against Manipulation

The combined effect of our decentralized components makes AnyDex resistant against a wide range of manipulations. We now outline how the mechanisms that make up AnyDex mitigate various types of manipulations. In this section, we first describe how TrustChain addresses a more overarching concern of electronic marketplaces, namely illegitimate modification of stored market information. Then we show how AnyDex addresses the manipulation concerns listed in Table 1. Finally, we describe in Section 5.9 how the XChange mechanism addresses counterparty fraud, the situation where a trader attempts to steal assets from a counterparty during a trade.

Modifying Market Information. AnyDex uses the TrustChain ledger to store market information. A key requirement is to detect fraud, i.e., whether a user has modified stored market information in their personal ledger after creation. This modification affects the integrity of TrustChain and would result in the situation where a user can hide prior records, therefore misleading other users. TrustChain addresses this as follows: a proposal or confirmation record is sent to a few random users after its creation. In addition, users periodically request records in the personal ledgers of random other users. When a user receives a record, the validity and consistency of the record is verified against the records that the user has already collected. The full description of this detection algorithm can be found in de Vos and Pouwelse (2020). This verification procedure is computationally efficient and merely takes milliseconds on average consumer hardware. Upon the detection of a modification, its evidence can be spread in

the network (e.g., the original and modified record) and other users can then refuse to communicate with the malicious user. Therefore, TrustChain mitigates the threat of a user manipulating information that already has been included on the ledger. This basic approach avoids the need to reach network-wide consensus on all records and results in superior throughput compared to existing distributed ledgers.

To quantify the speed of fraud detection, we have conducted various experiments in our prior work (de Vos and Pouwelse, 2020). We consider a network with 5'000 users where every user creates one proposal record per second. A record is sent to five random users after its creation, and every user requests two random records from two users every second. During the experiment, each user modifies a record in its personal ledger with 10% probability when creating a new record. Under these parameters, we find that on average, malicious modifications by users can be detected by other users within 3.6 s on average. 82.9% of all malicious modifications can be detected within five seconds after its occurrence. System designers can increase the intensity of record exchange to further reduce fraud detection times, at the cost of increased network traffic.

5.1. Front running

Front running occurs when a participant has priority access to market information. The distributed ledger used by AnyDex is designed such that participants have equal access to available information. Specifically, users send their TrustChain records to random other participants after creation, lowering the possibility for a participant to sustainably front-run on first-hand information. We should note that the threat of front running is reduced under the assumption that the underlying network graph is roughly k -regular, meaning that each user is connected to approximately the same number of other users. Would a particular user be connected to many other participants, it gains more market information compared to a user that maintains fewer connections. The network library used by AnyDex, however, addresses this situation by building a random overlay network.

Front running is also addressed by the MATCH mechanism. In particular, a new order is sent to a certain number of random matchmakers in the network. Additionally, an order creator waits some duration after receiving its first *match notification* before making a decision on which interested counterparty to pick for negotiation. This waiting window means that participants that learned about an order earlier do not have an advantage over participants that received the order slightly later.

5.2. Delaying or prioritizing orders

A particular threat in the MATCH mechanism are malicious matchmakers that deviate from the expected matchmaking behaviour, e.g., by prioritizing their own orders or by ignoring orders from particular traders. Since new orders in MATCH are sent to several matchmakers, it is likely that a suboptimal match suggested by a malicious matchmaker are superseded by ones from honest matchmakers. This approach makes the MATCH protocol resistant to malicious matchmakers that attempt to hide, prioritize, or delay particular orders.¹⁶ The protocol is highly scalable since it avoids the need to reach consensus on which orders are to be executed, unlike blockchain-based matchmaking approaches.

In our previous work, we have experimentally shown the resistance of MATCH against malicious matchmakers (de Vos et al., 2020). We highlight one of these experiments that focuses on the matching quality and fairness of MATCH in a ride-hailing environment as a decentralized alternative to Uber and Lyft. The workload is reconstructed from historical traces of taxi rides, comprising 2'100 ride orders and requests during 24 h in New York (TLC, 2017). We implement the matching policy such that it minimizes the distance between passengers and drivers, to reduce waiting times for passengers. Specially, the policy computes the geographic (haversine) distance between the locations included in orders and requests, and consider average distance between matched passengers and drivers as quality metric of a match. We model a malicious matchmaker as a driver that matches an incoming ride request from a passenger with its own service orders first. Intuitively, a new order should be sent to more matchmakers to counter the presence of malicious ones. Our experiment reveals that in a network with 2'000 matchmakers and 50% of all matchmakers being malicious, by sending an order to just 15 more matchmakers, the quality of matches in our mechanism is on par with the situation where all matchmakers are honest. Even when 75% of all drivers prioritize their own ride services during matchmaking, negotiated matches in our market maintain a high quality if we send new orders to 27 more matchmakers.

5.3. Price steering

Price steering occurs when different users receive different product results, or results in different order for the same query. Steering is possible because e-commerce operators are capable of using arbitrary metrics non-transparent metric like "Best Match" or "Most Relevant", rather than objective metric (price or user reviews).

In the AnyDex architecture price steering is mitigated by the decentralization of information management and matchmaking components. The order dissemination in MATCH is designed to ensure a high probability that a new order reaches an honest matchmaker with the current best matching order in their order book. Moreover, the specifications of matching policies (which may vary between different application contexts) are always transparent for the users. Users accumulate match notifications for a small period in their match queue (step 4 in Fig. 5). This provides a window for users to collect different matches from matchmakers. When sending a new order to a sufficient number of matchmakers, the quality of match notifications received from honest matchmakers will likely supersede those of malicious matchmakers. Then, a user starts to negotiate with the user that has the best counter-offer, and negotiates the next-best order if this negotiation is unsuccessful. As a result, no matchmaker can unilaterally force the execution of a particular counter order, even if it is the first to respond with a match notification.

¹⁶ Currently, MATCH is unable to identify and isolate a particular malicious matchmaker. This feature could be implemented by also recording the exchange of orders and match notifications on the TrustChain ledger, allowing any user to verify that a particular matchmaker correctly followed a matching policy. We consider this improvement beyond the scope of our work.

5.4. Price discrimination

Since information management in AnyDex is fully decentralized, our marketplace mitigates the threat of price discrimination. This means that a single participant is unable to influence what information other participants will receive from the rest of the network. We do note, however, that due to delays in the network, it is possible that two participants have differing views on the available offers for a particular product. This inconsistency, however, stems from the network layer and is not caused by a market operator controlling the flow of information. Participants can increase the rate at which market information is exchanged and fetched, reducing inconsistencies at the cost of increased resource usage.

5.5. Trader impersonation

A key type of manipulation in identity management is impersonation. Thus, the identity management component has to ensure unforgeability and Sybil-resistance. Unforgeability means that an adversary cannot forge the credentials of honest users or otherwise impersonate them. The Sybil attack manifests as a malicious user who attempts to subvert the system by joining the network under multiple identities (Douceur, 2002). Integration of the TCID system and using cryptographic keypairs with associated attestations addresses the threat of trader impersonation since this would require the private key of the trader being impersonated, unlike when identity data is managed by an authoritative market operator. Revocation of credentials can also help mitigate the threat of identity theft. SSI credentials data flows also remove the need for trusted third parties when a user authenticates itself to other users. A malicious user would only be able to delay the communication channel establishment, which is an attack targeted at the network layer and not at the market.

5.6. Censorship

Unbundling of identity management from information management and matchmaking in AnyDex provides high level of guarantees for censorship-resistance as no single party controls the process of traders onboarding to marketplace. Compared to identity owned by marketplace operator who can censor transactions based on traders' identity, in this setting only transacting parties participate in identity data flows, and issuers of credential have no knowledge to whom and when identity owner presents their credentials. In a decentralized marketplace environment, arbitrary censorship of traders would require collusion between the majority of users in other components such as matchmaking engine. However, as we demonstrate with matchmaking experiments even if 75% of matchmakers act maliciously they cannot prevent honest users from transacting.

5.7. Behaviour manipulation

In AnyDex, information is not served from a central point. This makes it infeasible for any single party to influence the behaviour of other users at a large scale. Hypothetically, the user interface (like any other interface) of AnyDex can be designed to steer the behaviour of users, e.g., by hiding or modified ranking of orders. At the same time, the specifications of the AnyDex protocol are public and developers have the ability to design alternative user interfaces. In the environment where backend trading engine is available through universally accessible API, providers of user interfaces are dis-incentivized to manipulate users, who can easily migrate to a different interface. Manipulation on the level of user interface is further obviated by the availability of aggregator services, providing users with search engines for prices and market offers. Empirical data shows that on markets for blockchain assets aggregators for decentralized exchanges (DEXs) occupy about 13.9% of market share (The Block Research, 2021).

5.8. Quality filtering

AnyDex provides several mechanisms to address manipulation of quality filtering. Integration of the Sybil-resistant TCID engine can enable a threshold for quality of user reviews since it is linked to their identity. For instance, it can guarantee that only reviews from marketplace users with a meaningful trade history are visible to other traders. More complex conditions for the quality filtering of offers can be implemented with a NetFlow component. The NetFlow algorithm is designed to provide accurate reflections on the honesty of traders, even if they try to inflate their own reputation or try to decrease that of others.

5.9. Counterparty fraud

Counterparty fraud, the situation where a counterparty refuses or is unable to fulfil its contractual obligations during a trade, is a major concern in electronic marketplaces (Ba et al., 2000). These risks are particularly pronounced in the context of a decentralized marketplace. We conducted an experiment and quantified how successful adversarial users are in committing fraud using XChange (de Vos et al., 2021). Our experiment shows that counterparty fraud does not present a systemic risk for AnyDex. For the experiment we use a real-world trading dataset, containing buy and sell orders published on the BitShares blockchain. We replay a week of trades in AnyDex, consisting of 125'527 buy orders, 104'423 sell orders and 212'489 cancellations of existing orders. These orders have been created by 1'161 unique users and involve 243 different assets. We consider the most challenging scenario where every user attempts to commit fraud. Without using the XChange mechanism, the total fraud gains of adversaries total to \$18.5 million. Using the XChange, adversaries are only able to gain \$16'260, a reduction of 99.9%. A key feature of XChange is that it bounds the total amount of fraud that a malicious user can commit. The TCID identity component prevents a malicious user from re-entering the market under a new identity after having committed fraud (also known as reputation whitewashing).

6. Related work on the components for decentralized marketplaces

The presented decentralized implementation of reference marketplace architecture demonstrates feasibility of technological separation. We do not suggest that some particular type of implementation is necessarily best suited for all types of markets and use contexts. We do argue that the backbone infrastructure for the open ecosystem of digital markets can be built on the basis of modular and interoperable basic components. The burgeoning growth of DeFi ecosystem provides an interesting vision of what such an ecosystem could look like in the future. Firstly, these are prominent examples of extreme disintermediation in financial services that were not thinkable even a few years ago. Secondly, the open nature of these solutions, both in terms of development and openness for participation, suggests strong potential for economic and technological experimentation.¹⁷ We argue that both of these factors also point out the future of decentralized electronic markets in general. It is important to note, of course, that current generation of DeFi protocols is exclusively focused on digital products, or so called on-chain assets. However, from the research perspective these solutions present bleeding-edge experiments in scalable decentralized infrastructures for marketplaces and as such deserve some closer scrutiny.

¹⁷ At the moment of this paper writing (Oct.2021) the value locked in DeFi protocols amounts to \$97 billions. See <https://defipulse.com/>

6.1. Information management

The GEM system, introduced already in 1999 by Reich et al. is one of the first published electronic markets where information is stored on different servers, spanning multiple geographic locations (Rachlevsky-Reich et al., 1999b). More recently, there have been various proposals to build electronic markets where different operators or participants themselves manage all information. In the PeerMart and OpenBazaar (Hausheer and Stiller, 2005; Arps and Christin, 2020).

However, the first practical decentralized electronic marketplace at scale was enabled by Distributed ledger technology (DLT). Blockchain based DEXes facilitate the direct exchange of assets between parties without central operator (Lin et al., 2019). All market information generated by traders is stored on and managed by self-enforced, contractual logic running on a blockchain, for example, smart contracts. Early experiments, such as EtherDelta, have revealed the major bottleneck for on-chain information management. Limited transaction throughput is not sufficient to facilitate the volume of major electronic markets that often process thousands of transactions per second (Zhou et al., 2020). An alternative approach to this problem that fuelled boom in DeFi, is a quote-driven alternative Automated Market Makers (AMM): smart contract that autonomously adjust the price for supply and demand based on incoming trade requests. One most interesting aspect of these solutions is that AMMs generate a lot of open-data, driving the development of third-party price aggregators (1inch, ParaSwap), that offer traders ability to find best prices. This emerging ecosystem of market information is very different from closed proprietary marketplaces where users have no access to pricing algorithms and other crucial market information.

6.2. Identity management

Current generation of DeFi solutions is bootstrapped by the minimal identity primitives provided by unique cryptographic keypair consisting of a public and secret key. Such keypair stored in a cryptocurrency wallet is de-facto all that trader needs to interact with on-chain protocols and DEXes, perform trades, borrow and issue loans. Such minimal identities, however do not allow for more complex functionality, which is why current DeFi crediting is mostly limited to overcollateralized loans.

There have been various proposals to introduce rich decentralized identity management solutions using DLT (Dunphy and Petitcolas, 2018). Self-sovereign identities (SSI) — one of the most promising and developed approaches — is a class of open standards and protocols for decentralized identity solutions providing end users have better control of their identifiers and credentials (Ferdous et al., 2019). This makes SSI schemas a radical departure from identity solutions provided by e-commerce platforms such as Facebook or Google (many electronic markets rely on these for on-boarding) or platform-specific identities such as those offered by Uber or AirBnb. SSI prevents various types of manipulations by market operators. Unlike proprietary identity solutions, SSI is an interoperable standard that prevents platform lock-in, reducing dependencies on trusted intermediaries for market participants.

6.3. Risk management

Collateralization is increasingly being used in asset marketplaces operating in the growing Decentralized Finance (DeFi) ecosystem, particularly when considering lending markets for digital assets (Werner et al., 2021). Within DeFi, risks management is a self-enforced process where the logic of a smart contract can slash the collateral of a fraudster and reimbursing wronged traders. The XClaim trading protocol, for example, relies on collateral deposits to enable asset trading between distinct blockchain ledgers (Zamyatin et al., 2019b). Other applications

use a semi-decentralized approach where a group of weakly trusted notaries is assigned to oversee trade and intervene when one of the parties does not adhere to the rules specified in the contract. The transparency aspect of blockchain benefits risk management because traders can make their own educated decisions about prospective counterparties by inspecting historical transactions. Also, fraud becomes more difficult to conduct since transactions are self-enforced and irreversible, under the premise that the underlying blockchain is secure.

6.4. Matchmaking

Except for the GEM market system, decentralized matchmaking is exclusively used in blockchain-based marketplaces, to the best knowledge of the authors. Most DEXes integrate the matchmaking process with blockchain logic. This process either relies on a smart contract to match known orders or executes the matchmaking logic as part of the transaction validation. Decentralized matchmaking on a blockchain fabric increases fairness since matchmaking proceeds according to predefined and self-enforced business rules that an authority cannot easily overrule. However, since users need to pay fees when creating transactions to manage their orders, order management can become costly. Furthermore, matching on a blockchain can be orders of magnitude slower than centralized matchmaking due to the need to reach a network-wide consensus on all transactions. Additionally, timing issues related to the consensus mechanism enables front running on orders (Eskandari et al., 2019).

Some blockchain-based marketplaces maintain the order book outside the blockchain to lower the costs of order management. Loopring, for example, is an order sharing protocol where new orders are sent to one or more relays in the network (Wang et al., 2018). Relayers claim the margin between two matched orders or can alternatively charge a fixed fee for their services. The Republic Protocol builds a decentralized network of nodes that match orders without revealing individual orders. The protocol uses cryptographic techniques to break down an order in multiple order fragments, which are distributed through the network, thus hiding the identity of the order creator.

6.5. Settlement

Notary-based schemes are an approach to settlement where approval by a group of credible notaries is required to settle a trade. This approach has seen increased adoption within the domain of blockchain-based trading. Notary schemes aim to partially alleviate the trust issues arising when relying on a single trusted intermediary by relying on the collective decision power of multiple entities. These notaries reach an agreement on the occurrence of particular events, e.g., that a trader has sent the promised goods to a counterparty. This approach can usually withstand the adversarial behaviour of a fraction of all notaries, and the damage caused by a single malicious notary is limited. The Interledger project is the most advanced approach in this direction (Thomas and Schwartz, 2015).

There have been experiments to conduct asset exchange without intermediary by leveraging cryptographic techniques, e.g., atomic swaps. The atomic swap is a coordination process that enables two parties to exchange blockchain-based assets (Herlihy, 2018). Atomic swaps eliminate the risk of losing assets to an adversarial trader or intermediate operator during the exchange.¹⁸ The main idea is that both parties during a trade lock their assets in a specialized transaction on

the blockchain so that no single party can claim both locked assets. This property is achieved with Hash-Timelock Contracts (HTLCs), a particular transaction type that leverages hash locks and time locks. A hash lock is a restriction that prevents the transfer of assets until a secret is revealed. A time lock is a primitive that prevents assets from being transferred until a specific time. This time lock prevents the assets being traded from being locked up indefinitely during an atomic swap. However, since one of the parties can decide whether to continue or abort the swap, it can effectively speculate on asset prices without a premium (Han et al., 2019).

6.6. Conflict management

In traditional marketplaces, conflicts are often resolved by the market operator, an intermediary or, when an intermediary cannot reliably determine the duped party, by a judge. However, the goal of many blockchain-based solutions is to devise riskless markets where, at least theoretically, disputes cannot occur. Nonetheless, some blockchain-based marketplaces rely on the honest behaviour of traders and require adequate methods to resolve conflicts, e.g., using escrow services (Goldfeder et al., 2017). This escrow may be a single entity, e.g., another user in the marketplace, or a group of users with some authority to resolve the dispute. The Bisq marketplace where users can trade cryptocurrencies has dedicated arbitrators that monitor the transactions during a particular trade and can recall exchanged funds through the use of multi-signature transactions (Beams and Karrer, 2017). Other solutions enabling data exchange between two parties leverage cryptographic techniques where one of the parties can submit a cryptographic proof to the blockchain to prove malicious behaviour of a counterparty (Dziembowski et al., 2018).

7. Discussion

We have presented a reference architecture and an implementation of an open decentralized marketplace, which is open, modular and interoperable. We have demonstrated that this approach can address issues of power abuse by marketplace operators acting as gatekeepers. From the long perspective, such an approach could provide a viable alternative to monopolistic marketplace platforms. The latest explosion of DeFi projects often overlapping with DEXes is also closely related to our work and provides some hints at the future of decentralized electronic marketplaces. We envision, however, the following three key obstacles that require future research and experimentation.

1. Interoperability. The unbundling of key components of electronic marketplaces introduces new research challenges, interoperability arguably being the most pressing issue. Interoperability is crucial for electronic marketplaces and is an essential ingredient that allows consumers to substitute one product with another that is manufactured by a different company (Choi and Whinston, 2000). With the proliferation of e-commerce platforms, interoperability issues are starting to be recognized, as is also evident by the European Commission's efforts to harmonize market communication and standards. The proliferation of distributed ledger technology has resulted in a vast and varying landscape of decentralized markets, many of them which are not interoperable (Zamyatin et al., 2019a). A few notable projects, like Polkadot and Cosmos, aim to build an "Internet-of-Blockchain", a large ecosystem comprising many interoperable blockchains (Siris et al., 2019). However, interoperability is still largely unsolved, and together with cross-chain communication remain the key open research issues.

2. Recentralization. Concurrent to the increasing popularity of decentralized solutions, we observe that some components of decentralized systems start to converge to a centralized structure. While it is hard to make accurate predictions in which direction the development of blockchain-based platforms will proceed, it is fair to say that scalability issues stand in tension with the issue of "recentralization".

¹⁸ Trading assets residing on a single blockchain can be achieved within the same transaction and is an atomic process, meaning that either all assets are exchanged, or nothing happens, under the premise that the blockchain is secure. Exchanging assets between different blockchain ecosystems is more involved and often brokered by a trusted intermediary that manages wallets in the involved ecosystems.

This tension is particularly evident in the context of the Ethereum blockchain, currently the most common layer to devise DEXes and other DeFi applications. Currently, Ethereum provides decentralization at the expense of significant transaction fees stemming from limitations of the underlying Proof-of-Work consensus algorithm. These fees present a prohibitively high barrier for many new market participants, creating a challenge for the scalability of such electronic markets. This, in turn, creates a risk that such markets could be captured by centralized parties like Binance chain or Facebook Diem (formerly known as Libra), providing lower transaction costs at the expense of decentralization.

Furthermore, some proposed solutions that increase blockchain performance also carry the aforementioned risks. Layer-two solutions where the bulk of transactions is processed outside the blockchain aim to lessen the load of the primary blockchain, introduces new economic incentives for participants. This has led to the concentration of network resources by a limited number of nodes (Lin et al., 2020). Similarly, Proof-of-Stake consensus has been proposed as a mean to increase transaction throughput while reducing transaction fees, but create risks of power concentration by affluent users (Fanti et al., 2018). While we do not suggest that these are insurmountable obstacles, we do suggest that the problem of technological and economic decentralization for blockchain-based markets and DLT technology, in general, is an open issue.

3. Software Security. Whereas trusted intermediaries are primarily responsible for reducing risks in traditional marketplaces, blockchain solutions are secured through participants' collective effort and cryptographic techniques. These platforms enable developers to deploy their preferred applications and define business logic in primitives like smart contracts. However, design weaknesses in both the underlying blockchain logic and decentralized applications, originating from economic and technical mechanisms, can result in significant economic losses. In 2016, a minor software vulnerability in the smart contract managing the Decentralized Autonomous Organization, or DAO, enabled an attacker to compromise about \$50 million worth of assets (Dhillon et al., 2017). Gudgeon et al. revealed an attack on Maker, one of the most significant DeFi applications in terms of market share, where an attacker would be able to steal \$0.5 billion worth of collateral with minimal effort (Gudgeon et al., 2020). And more recently, an attack on Ronin network resulted in a loss of approximately \$622 million (Hayward, 2022). Significant research effort must be spent to make these value systems technically and economically secure.

8. Conclusions

We have identified and presented technological solutions that can be implemented to create modular decentralized architectures for electronic marketplaces. We have presented a reference architecture and its implementation for all key functional elements: information management, identity management, risk management, matchmaking, settlement, and conflict management. We explored the problem of manipulation by marketplace platform operators and mapped different types of manipulation to the control of these key functional components. We argued that the logic of *technological separation* interpreted as decentralization of marketplace platform on architectural level, and at the level of individual functional components, can address identified manipulation issues. This approach is consistent with the ambitions of EU Digital Markets Act framework aiming to promote an open ecosystem of fair marketplace platforms. We argued that open ecosystems of interoperable and scalable components for electronic marketplaces could enable new types of peer-to-peer platforms not controlled by a single operator, which would make them much more resilient to manipulation practices. We observed a proliferation of experiments around distributed ledger technology, decentralized finance (DeFi), and other trust-less mechanisms for secure value exchange and trade. While these solutions may vary in technological maturity, they highlight a viable research direction towards decentralized architectures for general-purpose electronic markets.

CRedit authorship contribution statement

Martijn de Vos: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Georgy Ishmaev:** Conceptualization, Methodology, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Johan Pouwelse:** Conceptualization, Resources, Visualization, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Agmon Ben-Yehuda, O., Ben-Yehuda, M., Schuster, A., Tsafir, D., 2013. Deconstructing Amazon EC2 Spot Instance Pricing. *ACM Trans. Econ. Comput.* 1 (3), 1–20. <http://dx.doi.org/10.1145/2509413.2509416>, URL <https://dl.acm.org/doi/10.1145/2509413.2509416>.
- Alt, R., Zimmermann, H.-D., 2019. Electronic Markets on platform competition. *Electron. Mark.* 29 (2), 143–149. <http://dx.doi.org/10.1007/s12525-019-00353-y>, URL <http://link.springer.com/10.1007/s12525-019-00353-y>.
- Arps, J.E., Christin, N., 2020. Open market or ghost town? The curious case of OpenBazaar. In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 561–577.
- Azevedo, E.M., Weyl, E.G., 2016. Matching markets in the digital age. *Science* 352 (6289), 1056–1057. <http://dx.doi.org/10.1126/science.aaf7781>, URL <https://www.sciencemag.org/lookup/doi/10.1126/science.aaf7781>.
- Ba, S., Whinston, A.B., Zhang, H., 2000. The dynamics of the electronic market: An evolutionary game approach. *Inf. Syst. Front.* 2 (1), 31–40.
- Bamberger, K.A., Lobel, O., 2017. Platform market power. *Berkeley Tech. LJ* 32, 1051, Publisher: HeinOnline.
- Barwise, P., Watkins, L., 2018. The evolution of digital dominance. In: *Digital Dominance: the Power of Google, Amazon, Facebook, and Apple*. Publisher: Oxford University Press, pp. 21–49.
- Beams, C., Karrer, M., 2017. Phase zero: A plan for bootstrapping the Bisq DAO. URL: <https://docs.bisq.network/Dao/Phase-Zero.html>.
- Bokányi, E., Hannák, A., 2020. Understanding inequalities in ride-hailing services through simulations. *Sci. Rep.* 10 (1), 1–11.
- Bolton, G., Greiner, B., Ockenfels, A., 2018. Dispute Resolution or Escalation? The Strategic Gaming of Feedback Withdrawal Options in Online Markets. *Manage. Sci.* 64 (9), 4009–4031. <http://dx.doi.org/10.1287/mnsc.2017.2802>, URL <http://pubsonline.informs.org/doi/10.1287/mnsc.2017.2802>.
- Cabral, L., Haucap, J., Parker, G., Petropoulos, G., Valetti, T., Van Alstyne, M., 2021. The EU Digital Markets Act a Report from a Panel of Economic Experts. Tech. rep., European Union, Luxembourg: Publications Office of the European Union, p. 36, URL [doi:10.2760/139337](https://doi.org/10.2760/139337).
- Chen, Y., Bellavitis, C., 2020. Blockchain disruption and decentralized finance: The rise of decentralized business models. *J. Business Venturing Insights* 13, e00151.
- Chen, L., Mislove, A., Wilson, C., 2015. Peeking Beneath the Hood of Uber. In: *Proceedings of the 2015 Internet Measurement Conference*. ACM, Tokyo Japan, pp. 495–508. <http://dx.doi.org/10.1145/2815675.2815681>, URL <https://dl.acm.org/doi/10.1145/2815675.2815681>.
- Chen, L., Mislove, A., Wilson, C., 2016. An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace. In: *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, Montréal Québec Canada, pp. 1339–1349. <http://dx.doi.org/10.1145/2872427.2883089>, URL <https://dl.acm.org/doi/10.1145/2872427.2883089>.
- Choi, S.-Y., Whinston, A.B., 2000. Benefits and requirements for interoperability in the electronic marketplace. *Technol. Soc.* 22 (1), 33–44.
- Comission, E., 2020. Proposal for a regulation of the European parliament and of the council on contestable and fair markets in the digital sector (Digital Markets Act). COM(2020) 842 final. URL https://ec.europa.eu/info/sites/info/files/proposal-regulation-single-market-digital-services-digital-services-act_en.pdf.
- Dhillon, V., Metcalf, D., Hooper, M., 2017. The DAO hacked. In: *Blockchain Enabled Applications*. Springer, pp. 67–78.
- Douceur, J.R., 2002. The sybil attack. In: *International Workshop on Peer-To-Peer Systems*. Springer, pp. 251–260.
- Duggan, W., 2021. GameStop Congressional Hearing: Robinhood, Citadel Deny Collusion In Prepared Remarks. Yahoo!Finance URL <https://finance.yahoo.com/news/gamestop-congressional-hearing-robinhood-citadel-153030836.html?guccounter=1>.
- Dunphy, P., Petitcolas, F.A., 2018. A first look at identity management schemes on the blockchain. *IEEE Secur. Priv.* 16 (4), 20–29.

- Dziembowski, S., Eckey, L., Faust, S., 2018. Fairswap: How to fairly exchange digital goods. In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. pp. 967–984.
- EC, 2020a. Statement by Executive Vice-President Vestager on Statement of Objections to Amazon for the use of non-public independent seller data and second investigation into its e-commerce business practices. URL https://ec.europa.eu/commission/presscorner/detail/en/statement_20_2082.
- EC, 2020b. Statement by Executive Vice-President Vestager on the Commission proposal on new rules for digital platforms. URL https://ec.europa.eu/commission/presscorner/detail/en/statement_20_2450.
- Eskandari, S., Moosavi, S., Clark, J., 2019. Sok: Transparent dishonesty: front-running attacks on blockchain. In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 170–189.
- Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., Wang, G., 2018. Compounding of Wealth in Proof-of-Stake Cryptocurrencies. URL <http://arxiv.org/abs/1809.07468> [cs.LG].
- Ferdous, M.S., Chowdhury, F., Alsaifi, M.O., 2019. In search of self-sovereign identity leveraging blockchain technology. *IEEE Access* 7, 103059–103079.
- Fletcher, G.H., Sheth, H.A., Börner, K., 2004. Unstructured peer-to-peer networks: Topological properties and search performance. In: *International Workshop on Agents and P2P Computing*. Springer, pp. 14–27.
- Giaglis, G.M., Klein, S., O'Keefe, R.M., 2002. The role of intermediaries in electronic marketplaces: developing a contingency model. *Inf. Syst. J.* 12 (3), 231–246.
- Goldfeder, S., Bonneau, J., Gennaro, R., Narayanan, A., 2017. Escrow protocols for cryptocurrencies: How to buy physical goods using bitcoin. In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 321–339.
- Göldi, A., 2020. A blind spot for the dark side: the monopolies we didn't see coming. *Electron. Mark.* 1–2.
- Gudgeon, L., Perez, D., Harz, D., Livshits, B., Gervais, A., 2020. The decentralized financial crisis. In: *2020 Crypto Valley Conference on Blockchain Technology*. CVCBT, IEEE, pp. 1–15.
- Han, R., Lin, H., Yu, J., 2019. On the optionality and fairness of atomic swaps. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. pp. 62–75.
- Hausheer, D., Stiller, B., 2005. Peermart: The technology for a distributed auction-based market for peer-to-peer services. In: *IEEE International Conference on Communications*, 2005. ICC 2005. 2005. 3, IEEE, pp. 1583–1587.
- Hayward, A., 2022. Hacker Drains \$622M From Axie Infinity's Ronin Ethereum Sidechain. URL <https://decrypt.co/96322/hacker-622-million-axie-infinity-ronin-ethereum>.
- He, S., Hollenbeck, B., Proserpio, D., 2020. The Market for Fake Reviews. *SSRN Electron. J.* <http://dx.doi.org/10.2139/ssrn.3664992>, URL <https://www.ssrn.com/abstract=3664992>.
- Herlihy, M., 2018. Atomic cross-chain swaps. In: *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*. pp. 245–254.
- Jullien, B., Sand-Zantman, W., 2020. The Economics of Platforms: A Theory Guide for Competition Policy. CESifo Working Paper 8463, Center for Economic Studies and Ifo Institute (CESifo), Munich, URL <http://hdl.handle.net/10419/223535>.
- Jun, S., Rui, L., 2013. Manipulation prevention and hedging effectiveness: Optimal settlement window design for CSI 300 stock index futures. *Emerg. Mark. Finance and Trade* 49 (6), 52–66.
- Khan, L.M., 2018. Amazon—An Infrastructure Service and Its Challenge to Current Antitrust Law. In: *Digital Dominance: the Power of Google, Amazon, Facebook, and Apple*. 710, Oxford University Press, p. 805.
- Khan, L.M., 2019. The Separation of Platforms and Commerce. *Columbia Law Rev.* 119 (4), 973–1098. <http://dx.doi.org/10.2307/26632275>, URL <https://www.jstor.org/stable/26632275>.
- Lin, T.C., 2016. The new market manipulation. *Emory LJ* 66, 1253, Publisher: HeinOnline.
- Lin, L.X., Budish, E., Cong, L.W., He, Z., Bergquist, J.H., Panesir, M.S., Kelly, J., Lauer, M., Prinster, R., Zhang, S., et al., 2019. Deconstructing decentralized exchanges. *Stanford J. Blockchain Law & Policy* 2.
- Lin, J.-H., Primicerio, K., Squartini, T., Decker, C., Tessone, C.J., 2020. Lightning network: a second path towards centralisation of the Bitcoin economy. *New J. Phys.* 22 (8), 083022. <http://dx.doi.org/10.1088/1367-2630/aba062>, URL <https://iopscience.iop.org/article/10.1088/1367-2630/aba062>.
- Malone, T.W., Yates, J., Benjamin, R.I., 1987. Electronic markets and electronic hierarchies. *Commun. ACM* 30 (6), 484–497. <http://dx.doi.org/10.1145/214762.214766>, URL <https://dl.acm.org/doi/10.1145/214762.214766>.
- Markham, J.W., 1988. Front-running-insider trading under the commodity exchange act. *Cath. UL Rev.* 38, 69.
- Mathur, A., Acar, G., Friedman, M.J., Lucherini, E., Mayer, J., Chetty, M., Narayanan, A., 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proc. the ACM on Human-Computer Interact.* 3 (CSCW), 1–32. <http://dx.doi.org/10.1145/3359183>, URL <https://dl.acm.org/doi/10.1145/3359183>.
- Mavroudis, V., 2019. Market Manipulation as a Security Problem: Attacks and Defenses. In: *Proceedings of the 12th European Workshop on Systems Security - EuroSec '19*. ACM Press, Dresden, Germany, pp. 1–6. <http://dx.doi.org/10.1145/3301417.3312493>, URL <http://dl.acm.org/citation.cfm?doid=3301417.3312493>.
- Mavroudis, V., Melton, H., 2019. Libra: Fair order-matching for electronic financial exchanges. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. pp. 156–168.
- Mayzlin, D., Dover, Y., Chevalier, J., 2014. Promotional reviews: An empirical investigation of online review manipulation. *Amer. Econ. Rev.* 104 (8), 2421–2455.
- Menkveld, A.J., 2017. Crowded Positions: An Overlooked Systemic Risk for Central Clearing Parties*. *The Rev. Asset Pricing Stud.* 7 (2), 209–242. <http://dx.doi.org/10.1093/rapstu/rax016>, URL <https://academic.oup.com/raps/article/7/2/209/3806676>.
- Narayanan, A., Mathur, A., Chetty, M., Kshirsagar, M., 2020. Dark patterns: Past, present, and future: The evolution of tricky user interfaces. *Queue* 18 (2), 67–92.
- Otte, P., de Vos, M., Pouwelse, J., 2020. TrustChain: A Sybil-resistant scalable blockchain. *Future Gener. Comput. Syst.* 107, 770–780.
- Pavlov, V., Berman, R., 2019. Price Manipulation in Peer-to-Peer Markets and the Sharing Economy. Available At SSRN 3468447.
- Post, A., Shah, V., Mislove, A., 2011. Bazaar: Strengthening user reputations in online marketplaces. In: *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 11)*.
- Rachlevsky-Reich, B., Ben-Shaul, I., Chan, N.T., Lo, A.W., Poggio, T., 1999a. GEM: A global electronic market system. *Inf. Syst.* 24 (6), 495–518. [http://dx.doi.org/10.1016/S0306-4379\(99\)00029-0](http://dx.doi.org/10.1016/S0306-4379(99)00029-0), URL <https://linkinghub.elsevier.com/retrieve/pii/S0306437999000290>.
- Rachlevsky-Reich, B., Ben-Shaul, I., Chan, N.T., Lo, A.W., Poggio, T., 1999b. GEM: A global electronic market system. *Inf. Syst.* 24 (6), 495–518.
- Reich, B., Ben-Shaul, I., 1998. A componentized architecture for dynamic electronic markets. *ACM SIGMOD Rec.* 27 (4), 40–47. <http://dx.doi.org/10.1145/306101.306115>, URL <https://dl.acm.org/doi/10.1145/306101.306115>.
- Schechner, S., 2021. Facebook's marketplace faces antitrust probes in EU, U.K. URL <https://www.wsj.com/articles/eu-and-u-k-open-antitrust-probes-into-facebook-11622800304>.
- Schmid, B., Lindemann, M., 1998a. Elements of a reference model for electronic markets. In: *Proceedings of the Thirty-First Hawaii International Conference on System Sciences*. 4, IEEE Comput. Soc, Kohala Coast, HI, USA, pp. 193–201. <http://dx.doi.org/10.1109/HICSS.1998.655275>, URL <http://ieeexplore.ieee.org/document/655275/>.
- Schmid, B.F., Lindemann, M.A., 1998b. Elements of a reference model for electronic markets. In: *Proceedings of the Thirty-First Hawaii International Conference on System Sciences*. 4, IEEE, pp. 193–201.
- Schneider, T., Birkner, R., Vanbever, L., 2021. Snowcap: synthesizing network-wide configuration updates. In: *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. pp. 33–49.
- Scott, M., van Dorpe, S., 2020. EU charges amazon with misusing data, opens new probe over buy box. URL <https://www.politico.eu/article/amazon-antitrust-europe-margrethe-vestager/>.
- Siris, V.A., Nikander, P., Voulgaris, S., Fotiou, N., Lagutin, D., Polyzos, G.C., 2019. Interledger approaches. *IEEE Access* 7, 89948–89966.
- Stephens, E., Thompson, J.R., 2017. Information asymmetry and risk transfer markets. *J. Financial Intermediation* 32, 88–99. <http://dx.doi.org/10.1016/j.jfi.2017.05.003>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1042957317300360>.
- Stokkink, Q., Ishmaev, G., Epema, D., Pouwelse, J., 2021. A truly self-sovereign identity system. In: *2021 IEEE 46th Conference on Local Computer Networks*. LCN, IEEE, pp. 1–8.
- Subramanian, H., 2017a. Decentralized blockchain-based electronic marketplaces. *Commun. ACM* 61 (1), 78–84. <http://dx.doi.org/10.1145/3158333>, URL <https://dl.acm.org/doi/10.1145/3158333>.
- Subramanian, H., 2017b. Decentralized blockchain-based electronic marketplaces. *Commun. ACM* 61 (1), 78–84.
- The Block Research, 2021. Digital asset outlook 2022. The Block Research - GSR, p. 163, <https://www.theblockcrypto.com/post/127723/the-block-research-2021-digital-asset-outlook-report>.
- Thomas, S., Schwartz, E., 2015. A protocol for interledger payments. URL <https://interledger.org/interledger.pdf>.
- TLC, N., 2017. NYC taxi and limousine commission (tlc) trip record data.
- Trastour, D., Bartolini, C., Preist, C., 2003. Semantic Web support for the business-to-business e-commerce pre-contractual lifecycle. *Comput. Netw.* 42 (5), 661–673. [http://dx.doi.org/10.1016/S1389-1286\(03\)00229-9](http://dx.doi.org/10.1016/S1389-1286(03)00229-9), URL <https://linkinghub.elsevier.com/retrieve/pii/S1389128603002299>.
- Tzanetakis, M., Kamphausen, G., Werse, B., von Laufenberg, R., 2016. The transparency paradox. Building trust, resolving disputes and optimising logistics on conventional and online drugs markets. *Int. J. Drug Policy* 35, 58–68.
- Veit, D.J., Weinhardt, C., Muller, J.P., 2002. Multi-dimensional matchmaking for electronic markets. *Appl. Artif. Intell.* 16 (9–10), 853–869.
- de Vos, M., Ileri, C.U., Pouwelse, J., 2021. XChange: A universal mechanism for asset exchange between permissioned blockchains. *World Wide Web* 1–38.
- de Vos, M., Ishmaev, G., Pouwelse, J., 2020. MATCH: A decentralized middleware for fair matchmaking in peer-to-peer markets. In: *Proceedings of the 21st International Middleware Conference*. pp. 74–88.
- de Vos, M., Pouwelse, J., 2020. ?ConTrib: Universal and decentralized accounting in shared-resource systems. In: *Proceedings of the 1st International Workshop on Distributed Infrastructure for Common Good*. pp. 13–18.

- Wang, D., Zhou, J., Wang, A., Finestone, M., 2018. Loopring: A decentralized token exchange protocol. URL https://github.com/loopring/whitepaper/blob/master/En_whitepaper.Pdf.
- Werner, S.M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., Knottenbelt, W.J., 2021. SoK: Decentralized finance (DeFi). arXiv preprint [arXiv:2101.08778](https://arxiv.org/abs/2101.08778).
- Zamyatin, A., Al-Bassam, M., Zindros, D., Kokoris-Kogias, E., Moreno-Sanchez, P., Kiayias, A., Knottenbelt, W.J., 2019a. SoK: communication across distributed ledgers.
- Zamyatin, A., Harz, D., Lind, J., Panayiotou, P., Gervais, A., Knottenbelt, W., 2019b. Xclaim: Trustless, interoperable, cryptocurrency-backed assets. In: 2019 IEEE Symposium on Security and Privacy. SP, IEEE, pp. 193–210.
- Zhou, Q., Huang, H., Zheng, Z., Bian, J., 2020. Solutions to scalability of blockchain: A survey. *IEEE Access* 8, 16440–16455.