Mirqur.io: Pareto-optimal DEX on Cardano

lightpaper version 0.4

Chris M. Hiatt
lemmonade_714@protonmail.com

discord.gg/X4F8dVEJxG
twitter.com/Mirqurio
reddit.com/r/mirqurio
youtube.com/channel/UC6N6Z45hrneTXB2VBB_CD3w

Pool ticker: MRQR

July 5, 2021

Abstract

There have been many interesting innovations in the field of automated market maker (AMM) based decentralized exchanges (DEXes) over the past years. They suffered however from two major pathologies: Firstly, their contributions were separated over different products. Secondly, by running on the trailblazing but rather prototypical smart-contract-platform of Ethereum[1], they had either to constrain themselves from utilizing the full potential offered by financial mathematics or be constrained in practice by the enormous gas fees[2]. This paper outlines an approach to remedy both; the former by extensive and formally verified mathematical analysis and development, the latter by the choice of Cardano[3] as underlying smart contract platform.

1 Disclaimer

Note that this Lightpaper only serves as an outline; the detailed equations, their derivations as well as formal verification code will be released in time.

All of this might be subject to change, and any kind of critique or offer for collaboration are not only welcome but explicitly asked for. While no explicit scheme is defined for this herein, the intention is to reward everyone proportional to their contribution.

The reader is thanked for taking the time to read and apologies are offered in advance for any kind of grave mistakes that might as of now have

gone undetected. Any assistance through pointing them out is greatly appreciated.

2 What is "pareto-optimal DEX" supposed to mean?

"Pareto efficiency or Pareto optimality is a situation where no individual or preference criterion can be better off without making at least one individual or preference criterion worse off or without any loss thereof."[4]

In order to win, and continue winning, one either needs to deploy various tricks, or continuously hit some kind of optimum. The former would for example entail devising some scheme to lock in your customers, attack the competition on side channels or simply out-market everyone else.

Since that sounds rather tiresome however, Mirqur instead aims for two kinds of optima: The mathematical-economical optimum as well as the social-ethical one. In other words, no one shall be able to devise a more efficient offer for all participants, and no one shall shadow the inclusivity and fairness of the project. Note the partial ordering here: Of course other optima are desired - on the technical side, because there surely are some features we have not adopted as criteria at this point in time, warranting different DEXes; on the social side, because ideally, each project behaves optimally.

The latter part also means a project needs to be not only fair regarding contributors' compensation but designed with flexibility and updateability in mind from the get-go. It cannot be stressed enough that this can be your project too, if you want it to be. Being a miser towards people who helped bring something to life will hurt said thing; tokens will only be worth something if the attached project flies, and that will only happen with fairness. Questions of compensation will always be decided in concert with the community.

Hopefully the reader will excuse this small piece of moralizing now and enjoy the following technical feature introduction.

3 Preliminary concepts

This section briefly introduces some preliminary concepts. Feel free to skip.

3.1 AMMs and DEXes

A normal exchange works by actors publishing offers to buy or sell goods at a certain price; in the case of coincidence of wants, a trade is facilitated. This poses two problems:

1. The coincidence of wants needs to happen at the same time and place; there is no use selling Tomatoes in your backyard at 3AM in the morning, because (hopefully) there will be no one around wanting to buy them.
2. In case of price changes on other exchanges, the market maker (= actor awaiting the counterparty) needs to continuously update their prices. In a blockchain scenario, this comes with a pricetag. Were this not to happen very quickly, arbitrageurs could extract a lot of value from the market maker.

Automated market makers (AMMs) can be imagined as robots tasked with automatically adjusting the price based on balances; if for example your market stand contains twenty tomatoes and 100$, the price for one tomato would amount to

$$\frac{100\$}{20} = 5\$ \tag{1}$$

Then, after the successful liquidation of one tomato for the price of 5$, the new price would rise to be

$$\frac{100\$ + 5\$}{20 - 1} = \frac{105\$}{19} \approx 5.53\$ \tag{2}$$

per tomato. If, however, instead someone sells a tomato to your robot, the price for one fruit will adjust as follows:

$$\frac{100\$ - 5\$}{20 + 1} = \frac{95\$}{21} \approx 4.52\$ \tag{3}$$

This is the core functioning of a (simple) AMM[1]; it attempts to fit the price to current demand and supply by simply looking at its own ratios of goods and/or currencies.

The concept of a decentralized exchange (DEX) is nothing more than an exchange implemented as a smart contract on a blockchain platform (or analogous decentralization technology) instead of a centralized server. The obvious advantage is trust gained by the inability for either the creators of the system or the self-declared owners of the land the servers are physically located upon to compromise it; of course this is no guarantee, since backdoors can, will and have been implemented in many such projects. Therefore, publication of the code base is also crucial as is due diligence on part of the participants.

3.2 Cardano, Plutus and the EUTxO-model

Due to its decision to take a peer-reviewed, research-first approach as well as deploy formal verification to its code base, Cardano is by far the most trustable blockchain technology to date.

By using proof-of-stake and other innovations (in-built governance[5] and Ouroboros Hydra[6] come to mind) transaction costs are expected to

[1] To be more exact, the way Uniswap[12] v1 and v2 work

be kept low to a point of practical usability.

Plutus[7] - the native smart contract language - is based on Haskell[8], which, as a purely functional language offers a multitude of advantages for development, not the least of which are flexibility[2], natural formal verification and elimination of whole classes of bugs.

The extended-UTxO-model[9] on which Cardano operates - while being less intuitive than Ethereum's account-based model - lastly offers a range of advantages. Here is not the place to elaborate, but this fact has already proven clear during early development of Mirqur's smart contracts.

3.3 Slippage

Slippage[10] refers to the difference between expected and actual price in a trade. In the case of DEXes, this results from the way price is adjusted based on asset balances. Consider the tomato example again:

If the market stand contains 100$ and 20 tomatoes, the price for one tomato is $\frac{100\$}{20} = 5\$$. The same is true if the stand contains 10000$ and 2000 tomatoes - $\frac{10000\$}{2000} = 5\$$. The latter however would incur far less slippage - consider how the purchase of one tomato would adjust the price in each stand, and in turn how much one would have to pay for two tomatoes:

$$\frac{100\$ + 5\$}{20 - 1} = \frac{105\$}{19} \approx 5.53\$ \tag{4}$$

$$\frac{10000\$ + 5\$}{2000 - 1} = \frac{10005\$}{1999} \approx 5.005\$ < 5.53\$ \tag{5}$$

So, while two tomatoes in the less liquid stand would cost $5\$ + 5.53\$ = 10.53\$$ due to the purchaser buying a large percentage of the stock (5%), the latter, bigger stand would only ask $5\$ + 5.005\$ = 10.005\$$ for the two fruits. To put it differently - it incurs less slippage.

Therefore, it is important for an exchange's market makers to offer a lot of liquidity relative to trade size.[3]

3.4 Liquidity provision

The "pool" in liquidity pool comes from the fact that it pools the liquidity of multiple liquidity providers (LPs). The more liquid a market, the less slippage it incurs, and by pooling assets it is not longer required for cumbersome organizations to provide this service, which in turn obsoletes

[2] Recommended: http://www.paulgraham.com/avg.html

[3] In practice, not individual trade size but total trade size in one direction within a short timeframe is what counts. Consider two trades arriving at the smart contract within a very short timeframe; the latter one having only a very small tolerance in price. If the former changes the price too steeply, the latter order might fall flat, which is a rather annoying effect.

the overhead and various organizational slippages (cronyism, office politics and other waste resulting from unhealthy incentive structures come to mind).

In short, being an LP is a way for anyone with capital to act as a market maker and earn fees.

3.5 Impermanent loss

If an LP adds a number of assets to a pool which subsequently change their external valuation in relation to each other, arbitrageurs will trade with the pool until the price roughly reflects the external one. Were the LP now to withdraw their funds they would find that they would get less value out of it compared to simply holding the assets (not accounting for fees).

This loss is called "impermanent" since if or when the price returns to the initial ratio, the arbitrageurs will in turn act remove the loss. It can be shown that in the end it is the traders who pay the arbitrageurs for the service of "connecting" multiple exchanges, and they also gain from it, for the added liquidity reduces slippage.

The interested reader is being directed to the very intuitive explanation at [11].

3.6 Oracles

In order for a smart contract to process real-world data, the latter needs first to be written onto the blockchain in an unambiguous manner, for if one would just include an API-call in the code different miners might come to different results if they are delivered different data (i.e. because their API-calls happen at different points in time, or the data supplier provides different pieces of data to different IP ranges due to country or VPN-restrictions).

4 Tangential previous work

This section discusses some noteworthy innovations of the space which Mirqur chooses not to adapt for the forseeable future, for reasons as explained below.

4.1 Constant-product market makers (Uniswap)

Uniswap[12] is by far the most popular DEX on Ethereum. Rightfully so; their innovation was the insight that the constant value equation of CMMMs (constant-mean market markers) can in some cases simplified as follows:

$$V = \prod_i T_i^{w_i} \rightarrow V = \prod_i T_i \tag{6}$$

...if all w_i are the same, that is, if the ideal/normal state of the pool is one of equal distribution of all tokens included.

This was a tremendously useful innovation due to the otherwise prohibitively high transaction fees of Ethereum. On Cardano however there will likely be no need for such simplifications.

4.2 Curve

Curve (formerly Stableswap[13]) is specialized on trading stablecoins; therefore, they have in their equations a certain very flat range in their value-equation around ideal token ratios, in order to reduce slippage. The reasoning behind this is as follows - we need to have nonlinear value equations, since with linear ones (in the simplest case, constant-sum market makers: $V = \sum_i T_i$) the opportunities for arbitrage in the case of price differences to external exchanges would have no limit (unless one were to constantly adjust the price manually, which is a far from optimal or even feasible solution).

Since however stablecoins by design are expected to not change price a lot in relation to each other, having simple product-based value equations would result in a rather uneventful pool which in turn won't earn very noteworthy fees to the LPs. By flattening the value curve around the ideal price they essentially recreate for stablecoins a similar functioning that product-based equations display for currencies which are more volatile relative to each other.

Their innovation is brilliant. Mirqur instead is satisfied with improving capital efficiency between stablecoins with limit ranges as discussed above; while a far less impressive solution, the trade-off seems economical considering the added complications to the model and small number of stablecoins existing on Cardano at the time of writing (0). Of course this nevertheless would be a very interesting future route of research.

4.3 Kyber/DMM

Kyber[14] has an interesting model in which they adjust trading fees based on price movements, increasing them in more volatile markets. Their reasoning is that this is done also by real-life market makers too and protects them, and analogously is supposed to protect Kyber's liquidity providers.

While this is indeed a fascinating and elegant innovation it is not clear in how far this is supposed to provide protection - is a more volatile market not in the interest of the market maker or LP, for they earn from trades only? It seems rather to be an optimization, in the sense that the higher the demand for a product (in this case, the opportunity to trade) its price needs to be increased also to maximize profits.

Since Mirqur is not yet established and seeks to attract not only LPs but also traders and does not want to risk driving the latters' habits away,

such a scheme will not be part of its initial model. Note however that this is orthogonal to the equations discussed above; since it affects only the fees, it would not require fundamental changes to add it later, should the need arise.

4.4 Oracles

The purpose of oracles is to bring off-chain data onto a blockchain, such that it might be processed by smart contracts. In the realm of AMM the advantage provided is that instead of relying on arbitrageurs to adjust the price (and thus incurring impermanent loss) the pools can adjust their prices more intelligently.

They are - in the dectralized case - implemented by having a number of actors play a Shelling game about a certain data point[15][16]. In short, that means everyone has first to commit to their guess in a hidden fashion (i.e. by publishing a salted hash of it); in the next step, the players reveal it. This prevents dishonest actors from simply copying others' answers and thus avoiding doing the work themselves.

Then, the truth is compiled from all answers by some kind of averaging scheme. Everyone who guessed it right gets rewarded, everyone else punished. This works because coordinating is hard if the only way to coordinate with the other players is listening to the broadcast you know everyone else is also hearing and which says "we will meet at the truth".

This system will also need a dedicated token for that; if the stake the players need in order to participate would be currency with many other use cases, it would be easy to just buy sufficient majority and outvote everyone else; thus not only being able to manipulate the system as one sees fit but also getting paid for it. If however one needs to purchase a dedicated token first, then, after your attack had become publically known, people would refrain from using the oracles, thus devaluing the token you just bought the majority of. In other words, it's akin to buying a house in order to rob it.

Now the crucial weakness of this scheme is that if one intends to build an immense financial infrastructure on top of it, the stress placed on the system would also increase. If for example some financial contract would result in $n >> m$ units of value being gained if the oracle would report falsely, and the market capitalization of the oracle token would be below $2m$ (or in that ballpark, depending on the actual scheme used therein) it would become more and more tempting to just purchase the whole oracle, offsetting the value lost through devaluation of the staked tokens by the immense value gained through the manipulated derivative. If on the other hand the oracles were to required to hold all value equivalent to potential damages, the usage cost resulting from their opportunity cost would be immense.

Therefore Mirqur will abstain from using any oracles to adjust token

prices and prevent arbitrage. Surely it is possible to conceive a scheme to fix said issues, but at this point it seems more prudent not to rely on it.

5 Essential previous work and own contributions

In this section several previous innovations are introduced and the modifications as well additions Mirqur contributes to them put forth.

5.1 Portfolio liquidity provision

The liqudity pools offered by Uniswap[12] and Bancor[4][17] at the point of writing only offer very limited pools: They only may include two currencies and only at an equal ratio.

Balancer[18] excellently widened this by instead of using constant-product value functions (7) they use constant-mean ones (8), thus allowing for an adjustable ideal price point as well as liquidity pools containing multiple currencies. Mirqur builds upon this.

$$V = T_0 \cdot T_1 \tag{7}$$

$$V = \prod_i T_i^{w_i} \tag{8}$$

Here V constitutes the value to be kept constant by each trade (not accounting for fees collected), T_i the balances of the tokens and w_i their respective weights. For historical[5] and practical reasons it is commonly assumed that $\sum_i w_i = 1$. As can be seen in (8), a token with a higher weight excerts stronger influence onto V.

Note that in the case of $w_i = w_j \forall i, j$, (8) behaves exactly like

$$V = \prod_i T_i \tag{9}$$

which in the case of two tokens is exactly (7). Therefore, Balancer's method encapsulates the simpler one by Uniswap, at least mathematically (meaning: not accounting for differences in fees).

The advantages of this model are twofold: First, by allowing multiple currencies within the same pool, LPs willing to provide more than two

[4]It should be noted that Bancor also offers an impressive innovation in the form of single-sided liquidity pools; however, in order to protect against certain economical attacks they are subjected to multiple limitations, one of which is that they only apply to a number of whitelisted tokens. Mirqur strives to not having to rely on such a mechanism, and while an analogous method is being developed it is not ready enough to be announced as an upcoming feature.

[5]Those market makers are called "constant-mean-market-makers" (CMMMs) based on the weighted geometric mean: $\prod_i T_i^{w_i}$. This is not neccessary for AMMs; here only the ratios of w_i to each other (at least as long a trade involves exactly two tokens at a time).

currencies can do so within the same pool. While it is possible to model this with multiple dual-pools (7), the latter would require multiple trades routed over an intermediate currency. Depending on future developments it might even be prudent to consider this implementation.

Secondly, and most importantly, weights allow for a second way to influence the price besides token balances only. To see this, following [18], find that the marginal price for T_1 measured in T_0 is

$$p_{T_0,T_1} = \frac{\frac{T_0}{w_0}}{\frac{T_1}{w_1}} \tag{10}$$

In the simplified constant-product case (7) where $w_0 = w_1 = 1$, the price solely depends on T_0 and T_1. This conflicts with the LPs' needs to earn fees on a portfolio consisting of assets in proportion of their own choosing instead of as dictated by market conditions.

Further advantages exist; they are not to be discussed at this point.

5.2 Impermanent loss insurance

If the ratios of the tokens in a liquidity pool deviate from their assigned ideal (defined by their respective weights as described in subsection 5.1), the LP suffers so-called impermanent loss. The meaning of that is: They will lose value relative to just holding the assets in the defined ratio. It is called impermanent because were the ratios to return to their ideal state, the relative loss would also return to zero. This phenomenon is rather unfortunate, since the ratios of cryptoassets will seldom stay stable, and many investors expect certain singular assets they are knowledgeable about to outperform everything else. This scenario would then result in a loss for the LP, thus failing to service maximalists[6].

One ingenious solution proposed by Bancor[19] is therefore to provide insurance against said impermanent loss. The concrete implementation however suffers from two shortcomings: Firstly, the lost value is paid out in newly-minted network tokens, thus potentially creating a tragedy-of-the-commons situation[7]. Secondly, in order to prevent a certain type of

[6]The term commonly describes investors concentrating their liquidity onto a single asset in a category, convinced it will outperform the competition completely.

[7]The term constitutes a situation where the individuals' costs of using a limited shared resource is independent of actual usage, thus creating an incentive for everyone to drain said common resource. An example would be a river in which everyone can dump their trash at no personal cost (the river is considered large enough that it will be swept away quickly). Since however this logic applies to everyone irrespective of the others' actions, soon the river will become a trash heap.
In the case at hand the potential trouble could stem from too many LPs claiming their insurance in said token, which dilutes the supply, a cost in turn to be born by all token holders equally, irrespective of their claiming of the insurance. Note however that Bancor also has ways to reduce token supply, and one might argue that the potential loss to holders is offset by increased attraction of LPs. Still, at this point it seems uncertain how those factors will add up in the future.

malicious economic attack (not to be elaborated upon in here) the network only offers this insurance for a number of whitelisted tokens.

Mirqur's solution is as follows. Firstly, there will be dedicated and isolated funds for insuring against impermanent loss, to isolate the risk from the remaining protocol. Secondly, the maximum insurance amount will be limited in proportion to contributions to said fund (LPs can choose the percentage of earned fees they wish to contribute). Thirdly, the whitelisting will be replaced by having separate insurance pools for each token; thus the risk will be limited to pools expressing their trust in said token by including it.[8]

Mirqur already has a closed-form equation for impermanent loss in the flexible pools as defined in subsection 5.1.

The current intention is that the proceeds from the protocol staking the ADA in the MRQR-pool will form said insurance pool.

5.3 Smart tokens

A smart token, likewise an impressive innovation by Bancor[20] is a liquidty pool which does not hold the smart token but has the ability to forge and burn it. In their pure implementation that means the asset technically has an unlimited supply, although this is firstly limited by the exponentially rising price and can secondly be hard-capped by combination with range pool functionality (see subsection 5.6).

This offers two tremendous advantages. Firstly, it removes the need for the token seller to lock tremendous amounts of the paired token(s); Secondly, it allows capturing the long tail. To see the implications one has to image an artist wishing to fund their next work by pre-selling an according SFT (semi-fungible token). This would come with two problems: Firstly, the buyers would likely be stuck with it and unable to sell it, for there is little expectation to find coincidence of wants or even a platform facilitating this trade. Secondly, interested late-comers would be unable to purchase it for the same reasons. Smart token pools would however be able to facilitate trade between parties separated by time by constituting an always-present counterparty.

5.4 Novel token launch mechanism

Mirqur will offer a novel token launch mechanism. Without disclosing further details at this point, this will allow for everyone to hold their own "ICO" (quoted since the actual mechanism is not exactly what one would expect) with minimal hassle, directly followed by an automatic liquidity pool initialization. Of course the mechanism cannot effectively kept hidden after platform launch at the latest, so it appears acceptable to ask for

[8] The implications of this latter part are still to be determined in fullness; therefore it is subject to change at this point.

the readers' patience now.

If you are a developer looking for a platform to launch a token or NFT/SFT, please don't hesitate to reach out and share your needs.

5.5 On-chain oracles

It cannot be prevented that the data produced by an onchain-DEX will be utilized as oracles; why one would even consider attempting this is elaborated upon below in subsection 4.4. Therefore, Mirqur attempts to at least route those powerful market forces in a responsible direction by offering a number of running weighted averages for the convenience of the developer. By offering running weighted averages the temptation to build derivative contracts based on only the current price for time pressure's sake is counteracted; the latter would not only endanger confidence in all the technologies involved but incentivize market manipulation; again, as discussed below. By offering running averages as default interfaces however the aspiring market manipulator would have to fight the arbitrageurs over an extended period of time, which is (irrespective of Mirqur's own liquidity) a doomed endeavor, since one would have to move the entire crypto-market (and keep it there) for that same interval. If that were to succeed, the manipulation would be paid for handsomely and thus - as one might argue - well-deserved. Sarcasm aside, this is a failure case inherent to all unregulated financial markets and beyond the current scope of the project (Although it would indeed be tempting to try and solve the issue).

5.6 Range pools and diode pools

A range pool - the exciting innovation by Uniswap v3[21] - is a pool that will only be available as long as the assets' prices move between two predetermined ratios or price points. This has two advantages: Firstly, it increases capital efficiency. Secondly, it allows for the LP to constrain their potential impermanent loss and enforce instead of just encourage a certain ratio between their assets.

Additionally, Mirqur offers the innovation of diode pools: Herein trades are only allowed in one direction, but not the other. The intent is to service the needs of people looking to both provide liquidity but also leisurely accumulate one asset at the cost of another, less desired one, hoping to achieve a better price overall at the risk of partially missing liftoff.

5.7 Optimal trade distribution

A result of the high customizability of pools is a very fragmented environment. Therefore a method was developed to compute the optimal distribution of a desired trade over all available pools trading both tokens.

Note that this cannot be fully optimized before information about the exact fee structure is available; but given the EUTXO model and the overall low fees, the difference between theoretical and practical results is

expected to be a nonissue. Should it turn out to be one however a number of schemes offer themselves to address this.

6 Governance

The purpose of governance in this project is twofold:

1. determine spending of funds
2. greenlight smart contract updates

The general idea is that discussion happens beforehand and the vote is mainly used for legitimization.

While all token holders may vote on update proposals, in order to prevent proposal-spamming submission must be limited. In order to achieve this, not only is there a (very low) minimum requirement for submission, but each rejected proposal submits the tokens held to a cooldown period inversely proportional to their weighted votes (see subsection 6.1); During this period, those tokens may not be used to submit another proposal. This implies that a malicious actor splitting their tokens over many wallets will not only have to hold and lock many tokens but may have the option to flood the system only once in a very long interval.

However, in order for a potential opposition to be able to form, it is possible to delegate the votes without giving up control of the underlying tokens. In case of a rejection the cooldown will still apply to the delegated tokens.

6.1 weighted votes

In order to align incentives and prevent disinterested actors buying the token, affecting a vote, then selling right afterwards again as well as reward loyalty, each token's voting power will be (diminishingly) weighted by how long it has been held as well as how long its owner pre-commits to holding it. The latter consists of two parts - the rolling and the effective interval. The former can be changed at will; the latter is computed as follows:

$$e_e \leftarrow max(e_e, e_{c+r}) \tag{11}$$

With e_e being the end of the effective interval, e_c the current epoch and r the length of the rolling interval. In other words, any commitment will constantly be refreshed unless the rolling interval is reduced by the holder suffiently, at which point they naturally still have to wait for their previous lock to time out.

This brings the added benefit of being able to weight people higher who pre-commit to the project before knowing what the next vote will be about, thus disincentivising opportunistic voting behavior.

The token can not be sold before expiration of the effective interval.

7 Tokenomics

In total the project will have 1,000,000,000 tokens, to be distributed as follows:

1. 3% as founder's reward
2. 5% for future fund development, to be put under the protocol's control as soon as possible; although a small part of that might be spent for development beforehand (excluding current developer/SPO).
3. 12% under the protocol's control for undefined later use. For example - if governance should decide so - it could be used to have a protocol-owned sale initializing a liquidity pool as mentioned in subsection 5.4.
4. 6.9420% given to past stake pool delegators
5. 73.058% given to future stake pool delegators with a decreasing supply schedule.

Airdrops to pool delegators past and future will function as follows: Each epoch is assigned a fixed number of tokens. This number is then distributed proportionally over all delegators of that epoch, excluding the pool operator.

The tranche allocated to historical delegators are distributed equally over epochs; the tranche for future delegators will continually decrease according to the following scheme:

$$\frac{S \cdot Y}{100} \cdot \left(\frac{99}{100}\right)^{E-e_f} \tag{12}$$

Where E is the epoch in question, S is the total supply of one billion and Y the fraction of that assigned to future delegators (73.058%). e_f is the cutoff-epoch separating historical from future delegators (= launch date of the platform).

That implies that after three years (=219 epochs) ca. 90% of the last category will have been distributed:

$$\sum_{E=e_f}^{e_f+219} \frac{S \cdot Y}{100} \cdot \left(\frac{99}{100}\right)^{E-e_f} \tag{13}$$

$$= \frac{S \cdot Y}{100} \cdot \sum_{E=0}^{219} \left(\frac{99}{100}\right)^{E} \tag{14}$$

$$= \frac{S \cdot Y}{100} \cdot \left(100 - \frac{99^{1+219}}{100^{219}}\right) \tag{15}$$

$$\approx 0.89 \cdot S \cdot Y \tag{16}$$

Note that the corresponding series (and thus total supply) converges as follows:

$$\sum_{E=e_f}^{\infty} \frac{S \cdot Y}{100} \cdot \left(\frac{99}{100}\right)^{E-e_f} \tag{17}$$

$$= \frac{S \cdot Y}{100} \cdot \sum_{E=e_f}^{\infty} \left(\frac{99}{100}\right)^{E-e_f} \tag{18}$$

$$= \frac{S \cdot Y}{100} \cdot \frac{1}{1-\frac{99}{100}} \tag{19}$$

$$= \frac{S \cdot Y}{100} \cdot 100 \tag{20}$$

$$= S \cdot Y \tag{21}$$

Pool fees will stay low; they are currently set to 1%.

Also note that after platform launch the liquidity pools will delegate their ADA to MRQR too, thus earning tokens for LPs (which the "vanilla" delegators have to share with; hence the sudden spike in payout).

Lastly, in order to be eligible for consideration as historical staker, for those who delegated between Shelley-mainnet/MRQR-pool-launch (August 2020) and the Ides (15th) of March 2021 - when Mirqur was officially announced - potentially, a certain minimum loyalty is required; examination of the data will produce clear clusters. If you delegated after the 15.03.21 this filter will not apply to you in any case.

8 Future work

8.1 Cross-chain trading

By deploying wrapped tokens and other kinds of bridging technologies one will be able to trade with any DEX cross-chain, including Mirqur. It will however be vastly more convenient and thus friction-reducing to integrate such features directly. Since that seems a nontrivial task as well as orthogonal to the existing offer it will not be included in the initial design.

8.2 Order books

While we already technically include the traditional order book way of going about trading via diode range pools (see subsection 5.6), offering this as an explicit feature would possibly require a dedicated implementation, for efficiency reasons. Since this has not yet been investigated it is not (yet) part of the targeted initial feature set.

8.3 Monetary value of governance token

Should governance decide so, in the future, a small percentage of trading fees collected could be distributed among token holders. The reasoning behind that would be to align interest in future health of the system and

attach some external monetary value to the token, thus strengthening the decentralization mechanism - otherwise there would be a risk of easy accumulation of all governance tokens within one actor.

This however will not be part of the initial release for multiple reasons; the omnipresent potentiality of the community deciding to add it at any point in time will be assumed to serve as sufficient holding incentive until then.

8.4 Formal verification

All equations have been formally verified via the LEAN theorem proof assistant [22]. The proofs as well their code will of course be made public, latest at time of launch if not earlier.

8.5 Optimal trade routing

There still is the issue of optimal trade routing: Say there is a pool with tokens A and B and another one trading B against C, but none (or just a small one) trading A against C. Then, if one would want to trade A against C, it would be optimal to - depending on the scenario - to route a part or all of the trade in a way that first obtains B for the offered token and then, in a next step, to trade B for the desired currency.

While this is not yet part of Mirqur's final design, there are a number of ways to achieve this, some of them quite straightforward; it has and will be investigated further.

8.6 Detailed equations and community involvement

As soon as sufficient traction is reached to prevent appropriation by centralized forces the detailed equations, their proofs, the proofs' formal verification code and the DApp-codebase will be released.

The idea is to have this become the best DEX overall by virtue of attracting enough contributors, which in turn is intended to happen by virtue of meritocratic Anarchism[9]. This in turn demands however a certain amount of forethought, for it is a vulnerable state. As mentioned now almost ad nauseam, outreach and debate is actively asked for.

If this lofty outcome proves unattainable, satisfying consolidations would be to at least incentivize more people to improve the area's frontier of knowledge with this project and/or provide enough market pressure for other projects to adopt similar stances.

[9]To preempt any ideological misconceptions: What is meant is not chaos or lawlessness, simply the removal or fail-saving of central points of control as soon as feasible. Of course, in the beginning - especially before Goguen rolls out - community organization will be rather informal, but with less than 150 people our human hardware is already well-equipped to handle the various challenges.

9 Acknowledgements

The author would like to thank everyone who helped bring Ethereum and Cardano to life, as well as the researchers cited directly and indirectly and everyone else who served the space thus far. Special gratitude is extended to his delegators, frens as well as all the other folks showing interest and support.

References

[1] https://ethereum.org/en/

[2] https://crypto.com/defi/dashboard/gas-fees

[3] https://cardano.org/

[4] https://en.wikipedia.org/wiki/Pareto_efficiency

[5] https://cardano.org/governance/

[6] https://iohk.io/en/blog/posts/2020/03/23/from-classic-to-hydra-the-implementations-of-ouroboros-explained/

[7] https://docs.cardano.org/projects/plutus/en/latest/

[8] https://www.haskell.org/

[9] https://jmchapman.io/papers/eutxo.pdf

[10] https://www.investopedia.com/terms/s/slippage.asp

[11] https://finematics.com/impermanent-loss-explained/

[12] https://uniswap.org/

[13] https://curve.fi/files/stableswap-paper.pdf

[14] https://files.kyber.network/DMM-Feb21.pdf

[15] https://chain.link/whitepaper

[16] https://github.com/Emurgo/Emurgo-Research/blob/master/oracles/Oracle-Pools.md

[17] https://bancor.network/

[18] https://balancer.fi/whitepaper.pdf

[19] https://drive.google.com/file/d/1en044m2wchn85aQBcoVx2elmxEYd5kEA/view

[20] https://storage.googleapis.com/website-bancor/2018/04/01ba8253-bancor_protocol_whitepaper_en.pdf

[21] https://uniswap.org/whitepaper-v3.pdf

[22] https://leanprover.github.io/