

Cellular structure for a digital fiat currency

Robleh Ali

MIT Digital Currency Initiative

Abstract

Bitcoin introduced digital money without banks. It showed that functions of the financial system could be reliably executed by decentralized networks and in doing so raised the prospect of a new financial system. Almost a decade later, Bitcoin still works despite frequent predictions of its demise but the overwhelming majority of payments still use fiat currencies and the financial system is unchanged. We propose a cellular structure in which digital fiat currency issuers and transaction validators create functionally separate ledgers, decentralized but capable of cohering into a single system for moving digital money. The primary benefit of the cellular structure is that it lowers barriers to entry for payments by using trustless intermediation between cells in the system. The larger purpose of this structure is to create an open foundation for a decentralized financial system in which competition can thrive but which cannot be captured by private interests.

1 Introduction

This paper addresses how digital fiat currency (DFC) could be provided to the financial system in the future. The goal of doing so is a changed organizational structure for payments with greater competition by reducing barriers to entry for participants in the system.

Existing payment systems usually comprise an operator which is responsible for running the system and members who participate in it. Together, the operator and members provide services to end users. Large payment systems tend to be centralized, the practical effect is that control over payment systems is concentrated making it difficult for new entrants to compete. This problem has manifest itself across different jurisdictions, leading to antitrust cases against payment system operators.¹

¹See, for example the European Union case against Visa and the US Justice Department cases against Visa and American Express.

1.1 Trust and its effect on structure

The core problem facing any digital payment system is how to prevent double spending. This is done with rules defining how the system should work and enforcement mechanisms to ensure the rules are followed. How rules are constructed and enforced affects the structure of the system. In existing payment systems, rules are extrinsic to the technology meaning that a combination of contractual obligations and regulations are needed to maintain the integrity of the system. The mechanisms for ensuring compliance with these rules are (a) new members need permission to join the system and (b) once inside the system there are penalties if the rules are not followed. Ultimately, these systems require trusted participants to operate, the rules are there to ensure they are worthy of that trust.

Bitcoin demonstrated it was possible to solve the double spend problem without trusted participants. It achieves this by making the rules and enforcement intrinsic to the protocol and technology. Participants do not need to be trusted because if they don't follow the rules, their transactions will be rejected. The only requirement for joining the system is running the software necessary to conform to the protocol - there are no requirements beyond that. The absence of pre-requisites flowing from extrinsic rules reduces the barriers to entry for the system.

It is true that the barriers to entry for Bitcoin mining have steadily increased over time through rising infrastructure costs. This has led to centralization of the mining network because small scale mining operations are now uneconomic. There is debate about the extent to which this is problematic for Bitcoin but it is the current reality. It does demonstrate that decentralized architecture alone is insufficient to prevent increased barriers to entry. In spite of this, Bitcoin is largely an open system, it is still possible to participate in the system (for example by running a full node) and build services on top of the platform without anyone's permission.

Generally the more trust is needed for a system to function, the higher barriers to entry will be. The DFC structure we propose is different from Bitcoin because a DFC is not the sole manifestation of the currency and requires backing as a result. When proposing the design for the system we do make use of the general concept of reducing trust in a system leading to lower barriers to entry overall.

2 Cellular structure

The structure we propose is a system built up from multiple different ledgers recording transactions for the same currency. Each ledger, or cell, in the system would have an issuer - a private entity which holds central bank money equal to the amount of DFC in issue on their cell. Routing payments between cells in the system would be done trustlessly.

2.1 Backing

The most significant difference between existing cryptocurrencies and DFC is backing. Cryptocurrencies have no backing and the value of the currency is set by the market whereas in DFC the issuer of the money on any given cell has to maintain 1:1 backing. How this backing works also affects the structure of the system. The goal is to minimise the underlying risks of the backing asset such that the regulatory structure can be kept simple. A characteristic of the existing financial system is its complexity which makes it difficult to manage risks for participants and regulators alike - see Haldane and Madouros (2012)[7]. Complexity makes regulation more costly to comply with and raises barriers to entry.

In our model, each cell in the system is backed by central bank money. This model of issuers holding cash at the central bank to back the DFC is a simplified version of the requirements placed on issuers of private banknotes in the UK.² In the UK system the assets held at the central bank are ringfenced and this allows all the privately issued notes to circulate at par with central bank issued banknotes. It is possible to back a DFC with alternative assets but these introduce either credit or market risk into the system, making it less stable and more complex. We have not addressed these potential models of a DFC system because the risk they introduce would require additional complexity to manage with no obvious benefit.

2.2 Single cell

A single DFC cell is comprised of an issuer and validators. In addition to providing the backing, the issuer is responsible for the software and protocol used to make transactions within the cell. The validators record transactions and maintain up to date copies of the ledger.

²For more information, see Bank of England: Scottish and Northern Ireland banknotes.

Figure 1: Single DFC cell




Figure 1 shows a decentralized model in which the validators are separate entities from the issuer and in which validators use distributed consensus to synchronize their ledgers. It would also be possible to have an internal construction for a cell which was more centralized where the issuer also undertook the validation of transactions and maintenance of the ledger. Our DFC implementation will use a Cryptokernel³ based blockchain with validators separate from the issuer.

Regardless of the internal structure of an individual DFC cell, the presence of an issuer does create a degree of centralization (as compared to a permissionless cryptocurrency). Having multiple DFC cells ensures that the system does not become reliant on a single issuer. All the DFC cells must support Hashed Timelock Contracts to enable trustless intermediation between the cells.

2.3 One currency - multiple cells

Having multiple cells allows for trustless routing of payments between them. This allows the system to reduce the barriers to entry for intermediaries as the rules are intrinsic to the technology, in this case Hashed Timelock Contracts (HTLCs).

³Cryptokernel is a blockchain toolkit developed by James Lovejoy at the MIT Digital Currency Initiative.

Figure 2: DFC system structure – trustless routing using HTLCs





Figure 2 shows a payment from A to B via C, a trustless intermediary, using HTLCs. This example assumes that there are pre-existing payment channels between both A and C as well as B and C (for a description of how payment channels work see Poon and Dryja (2016)[10]).

In essence, HTLCs work by B creating the hash of a secret "R" (HashR) which is then used by A and C to create HTLCs, one sending money from A to C and one sending from C to B. The money held in the HTLCs can be claimed by revealing R. A and C can compare HashR to verify that receiving R will allow both HTLCs in the path to complete. B triggers the payment from C by revealing R to C. This allows C to trigger the payment from A to C and complete the transaction. Neither A nor B need to trust C because she can only claim the funds from A after she has paid to B. C does not need to trust A because she knows the secret she receives from B will trigger the HTLC with A. This allows intermediation between cells to be trustless. If B never reveals R, then A and C can cancel the payment after a timeout and retrieve their funds, creating a strong incentive for B to reveal R immediately.

Figure 3: DFC system structure – one currency



Under this structure there are two ways for holding balances. One is for an end user to hold all their funds on a single cell and rely on intermediaries to

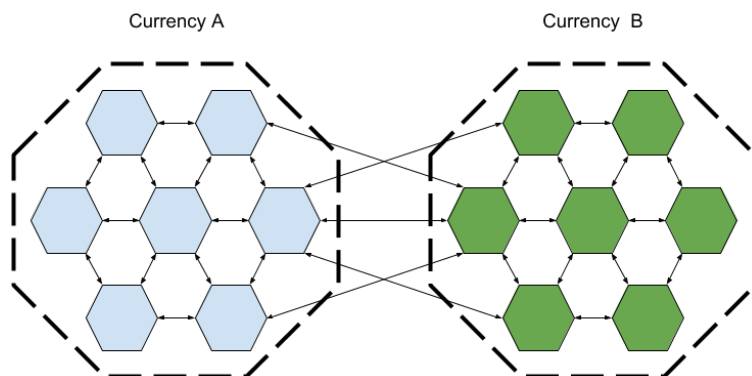
route money when they need to make payments to other users holding money on different cells in the system. In this model, risk arises from users being exposed to the operational failure of the cell they are using to hold their funds. This risk is similar to the current financial system in which users generally hold their funds at one bank.

An alternative is for users to hold balances across multiple cells. This increases a user’s risk of being exposed to failure of a given cell but the consequences are less severe as they would still be able to access their balances on other cells while the faulty cell is recovered.

In both models intermediaries would be required to route payments between different cells. Intermediaries in the cellular structure essentially perform the same function payment systems do today by connecting the ledgers of commercial banks, albeit in a different way. The difference is that intermediaries in a cellular system do not have to manage credit and liquidity risk in the way existing payment systems do, just the payment channels needed to route the money.

2.4 Multiple currencies

Figure 4: DFC system structure – multiple currencies



One advantage of the cellular structure is that the same infrastructure used to route payments between cells of the same currency can also be used to make transactions between different currencies. In addition to routing the payments between different cells, intermediaries would also have to make markets for the currency pairs users wanted to transact. Whether the same intermediaries would do both the market making and transaction settlement is an open question, the functions could be separated but this is not necessary and there may be benefits to allowing the same intermediary to do both.

3 Rationale for DFC with a cellular structure

The most commonly asked question when discussing DFC is why bother at all - payment systems to transact these currencies already exist and function adequately. The short answer is to add competition and innovation while improving the stability and resilience of a financial system whose serious flaws were exposed during the 2008 financial crisis. While it is true that the crisis did not emanate from payment systems, fundamentally reforming the financial system has to start with its foundation which is money and payments.

3.1 Competition

The cellular structure increases competition in the system by lowering barriers to entry. It achieves this in two ways. One is at the individual cell level, the relative simplicity of the regulatory structure means that the barriers to adding a cell are as low as possible. The other is by making routing between cells trustless, this also minimises barriers to entry for intermediation. The cellular structure accommodates competition from multiple private participants at both the infrastructure and intermediation levels which makes it more difficult for a single entity to capture the platform.

Payment systems are the foundation of the financial system. The goal of creating a DFC is an infrastructure model where new entrants have equal access without the permission of their competitors. This is one way in which the cellular structure supports competition within the DFC system. It also provides competitive pressure to existing payment providers, giving merchants and their customers a viable alternative. This is especially important as cash usage declines for day to day payments.

Equally important is ensuring the future infrastructure is not itself captured by any commercial interest which can then extract rents from use of the system and control who has access. It is absolutely critical that any software employed for transacting in a future financial system is not dependent on any cryptocurrency, current or future.

Many existing cryptocurrencies have inflated values predicated on the notion that many or all future financial transactions will flow across their networks, increasing demand for their tokens. The prospect of the financial system becoming reliant on a single cryptocurrency or small handful of cryptocurrencies would be a disastrous outcome. Putting aside the issue of whether these networks will ever be capable of processing the number of transactions needed to support the entire financial system, the prospect of access to the financial system requiring the use of a token whose ownership is highly concentrated is the antithesis of decentralization, whatever the underlying architecture of the system. Changing

the technology but keeping the same organizational structure (or making it even worse) is an exercise in futility.

3.2 Innovation

There are two different types of innovation to consider, technological and organizational. Technological innovation enables organizational innovation but doesn't guarantee it. In chapter 7 of their book, Brynjofsson and McAfee (2014)[2] give the example of electrification of factories (technological innovation) not leading to any significant productivity for more than thirty years. The change only came after a generation of factory managers had retired and were replaced by a new generation who saw that electric dynamos could be shrunk, allowing machines to be individually powered and factories remodeled around the flow of materials (organizational innovation). David (1990)[3] found that it was these organizational changes that the new technology enabled were primarily responsible for the productivity gains - not the mere introduction of new technology.

There is an analogy to the financial system and the emergence of cryptocurrencies. Digitizing money and other financial assets (technological innovation) happened as soon as computers were deployed in the financial system but there was no accompanying organizational innovation so the system didn't change. The current model of the financial system is for institutions to handle trust related aspects of the system and for software to handle record keeping. The change Bitcoin introduced was to handle both the trust element and record keeping using a combination of software and cryptographic methods. Trust in this context means two things (a) the integrity of the record keeping (b) the conduct of monetary policy.⁴

It was Bitcoin that first demonstrated that creating digital money without banks (organizational innovation) was possible. One goal of creating a DFC with a cellular structure is to lay the foundation for organizational innovation in the broader financial system. It would be perfectly possible to implement a DFC without any accompanying organizational change but doing this would be akin to factory managers' first use of electricity and the result would likely be the same - no appreciable gain in productivity.

Subsequent innovations in cryptocurrencies such as Hashed Timelock Contracts (HTLCs) and smart contracts⁵ have demonstrated how other functions of the financial system (payment versus payment, delivery versus payment, derivatives) could be handled using software rather than institutions. For example, in

⁴Monetary policy could in theory be handled automatically but this would require a more fundamental set of changes to the way money is created in the economy and is beyond the scope of this paper, for a discussion see Friedman (1960)[5].

⁵On Bitcoin for example, see Discreet Log Contracts by Thaddeus Dryja at the MIT Digital Currency Initiative.

a future where both DFC and tokenized shares exist, HTLCs eliminate the need for a trusted intermediary as the software ensures simultaneous transfer of the cash and asset - enabling delivery versus payment. The same method can be used to exchange two DFCs trustlessly to facilitate payment versus payment. This has systemic implications as the market for intermediation will become much more competitive as barriers to entry are lower. Any intermediary overcharging for services could be immediately routed around by users. This is a significant difference with the existing system in which the costs of intermediating payments means setting up a new payment system and getting all the existing banks to use it - a very high bar to competition.

Smart contracts allow two parties making a prediction on a future event such as the price of an asset at a specified date with the settlement enforced using software. Smart contracts have to be fully collateralized and are not completely trustless as they require an oracle to determine price of the asset on the execution date. A risk is that the oracle is compromised, one way of mitigating this risk is to use multiple oracles and use some subset of them to settle the transaction but this does not eliminate the risk completely. At present, smart contracts still have limitations which would prevent them from replicating the complete functions of existing derivatives. For example, because smart contracts are pre-funded on both sides, derivatives using DLCs only work in a relatively small price window (unless the counterparties are prepared to pre-fund with increasingly large amounts). There is also no netting which would make using smart contracts at scale in their current form very capital intensive. Smart contracts are the subject of ongoing research to address these limitations and the technology will continue to develop. They are also an example of an innovation from cryptocurrencies with the potential to alter the organizational structure of the financial system by using software to perform functions previously undertaken by institutions.

3.3 Financial stability and resilience

The primary way in which a DFC increases financial stability is by creating a payment system separated from credit risk. This reduces the risk of a macroeconomic shock from the simultaneous failure of the banking system and payments in the economy, reducing one incentive for the state to bailout the banking system.

Partially separating payment and credit in this way would have costs. The advantage of fusing deposit taking and lending is that it allows more efficient liquidity management such that money used for payments doesn't lie inactive on commercial banks' balance sheets, it can simultaneously support productive enterprise, see Kashyap, Rajan and Stein (2002)[8]. The cost of this efficiency is structural fragility which is mitigated but not eliminated by deposit insurance

and regulation. These measures introduce government subsidy and barriers to entry to the financial system which alter incentives and create costs of their own. Depositors no longer need to scrutinize the safety of banks because their deposits are government guaranteed. Banks can take greater risks, reduce capital and increase leverage in the knowledge that the losses will ultimately be borne by the taxpayer. In theory, regulation addresses this moral hazard problem but experience shows financial institutions will both push to the limit of what regulation allows and lobby for looser regulation. In addition, risky activity is shifted outside the regulatory perimeter and the financial crisis a decade ago demonstrated how risks in the shadow banking system can spill into the rest of the system - see Gorton and Metrick (2010)[6].

As the DFC in our model is fully backed by central bank money, there is a risk that in a crisis bank depositors could run to the DFC as a perceived lower risk form of money. Broadbent (2016)[1] describes this risk in the context of a central bank issued digital currency and also the potential effects on lending. As a DFC is a private system and deposit insurance would still exist these risks would be lower but the critique is still relevant to a DFC backed with central bank money. The risks could be mitigated by adopting the 'pawbroker for all seasons' proposal in King (2016)[9] which has the additional benefit of increasing the discipline on commercial banks. Shifting funding of productive enterprise from banks to capital markets would also reduce the systemic reliance on commercial banks - see Véron and Wolff (2016)[11] and Draghi (2014)[4].

The cellular structure increases resilience by ensuring the system can still function and route payments even if one or more of the cells isn't functioning. This structure also allows for multiple issuers, reducing the probability of any individual provider dominating the system.

4 Conclusion

Bitcoin showed that digital money could be separated from banks. It also demonstrated that functions of the financial system could be executed trustlessly by a decentralized network without the current set of institutions. Although Bitcoin has shown remarkable resilience since it first emerged almost a decade ago, it has not evolved into a platform capable of replicating all the functions of the existing financial system either in terms of complexity or scale.

Our work on DFC is aimed at progressing the concept introduced by Bitcoin - that of a decentralized financial system which is easy to access but hard to capture. There are levels to it. The cellular structure we have outlined in this paper shows how a DFC can be constructed in a decentralized way to preserve the possibility of organizational innovation in the levels above. This is important because money is the base level of any financial system and its structure affects

how the rest of the financial system built on top develops.

References

- [1] Ben Broadbent. “Central banks and digital currencies”. In: *speech at the London School of Economics 2* (2016).
- [2] Erik Brynjolfsson and Andrew McAfee. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company, 2014.
- [3] Paul A David. “The dynamo and the computer: an historical perspective on the modern productivity paradox”. In: *The American Economic Review* 80.2 (1990), pp. 355–361.
- [4] Mario Draghi. “Keynote speech at the Eurofi Financial Forum in Milan”. In: *European Central Bank, September, available at https://www.ecb.europa.eu/press/key/date/2014/html/sp140911_1_en.html* (2014).
- [5] Milton Friedman. *A program for monetary stability*. Vol. 541. Fordham University Press New York, 1960.
- [6] Gary Gorton and Andrew Metrick. “Regulating the shadow banking system”. In: *Brookings papers on economic activity* 2010.2 (2010), pp. 261–297.
- [7] Andrew Haldane and V Madouros. “The dog and the Frisbee. Bank of England”. In: *Speech given at the Federal Reserve Bank of Kansas City’s 36th economic policy symposium, “The Changing Policy Landscape”, Jackson Hole, Wyoming*. Vol. 31. 2012.
- [8] Anil K Kashyap, Raghuram Rajan, and Jeremy C Stein. “Banks as liquidity providers: An explanation for the coexistence of lending and deposit-taking”. In: *The Journal of Finance* 57.1 (2002), pp. 33–73.
- [9] Mervyn King. *The end of alchemy: Money, banking, and the future of the global economy*. WW Norton & Company, 2016.
- [10] Joseph Poon and Thaddeus Dryja. “The bitcoin lightning network: Scalable off-chain instant payments”. In: ().
- [11] Nicolas Véron and Guntram B Wolff. “Capital Markets Union: a vision for the long term”. In: *Journal of Financial Regulation* 2.1 (2016), pp. 130–153.