zkFUND:

Secure Decentralised Generalised One-time Ring Signature Peer-to-Peer Scalable Off-Chain Untraceable Electronic Instant Cash System and MimbleWimble Transaction Ledger Consensus Algorithm

ESPAcoin team

Abstract

A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution. Digital signatures provide part of the solution, but the main benefits are lost if a trusted third party is still required to prevent double-spending. We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort basis, and nodes can leave and rejoin the network at will, accepting the longest proof-of-work chain as proof of what happened while they were gone.

While several consensus algorithms exist for the Byzantine Generals Problem, specifically as it pertains to distributed payment systems, many suffer from high latency induced by the requirement that all nodes within the network communicate synchronously. In this work, we present a novel consensus algorithm that circumvents this requirement by utilizing collectively-trusted subnetworks within the larger network. We show that the "trust" required of these subnetworks is in fact minimal and can be further reduced with principled choice of the member nodes. In addition, we show that minimal connectivity is required to maintain agreement throughout the whole network. The result is a low-latency consensus algorithm which still maintains robustness in the face of Byzantine failures. We present this algorithm in its embodiment in the zkFUND Protocol.

The intent of zkFUND is to create an alternative protocol for building decentralized applications, providing a different set of tradeoffs that we believe will be very useful for a large class of decentralized applications, with particular emphasis on situations where rapid development time, security for small and rarely used applications, and the ability of different applications to very efficiently interact, are important. zkFUND does this by building what is essentially the ultimate abstract foundational layer: a blockchain with a built-in Turing-complete programming language, allowing anyone to write smart contracts and decentralized

applications where they can create their own arbitrary rules for ownership, transaction formats and state transition functions. A bare-bones version of Namecoin can be written in two lines of code, and other protocols like currencies and reputation systems can be built in under twenty. Smart contracts, cryptographic "boxes" that contain value and only unlock it if certain conditions are met, can also be built on top of the platform, with vastly more power than that offered by Bitcoin scripting because of the added powers of Turing-completeness, value-awareness, blockchain-awareness and state.

The bitcoin protocol can encompass the global financial transaction volume in all electronic payment systems today, without a single custodial third party holding funds or requiring participants to have anything more than a computer using a broadband connection. A decentralized system is proposed whereby transactions are sent over a network of micropayment channels (a.k.a. payment channels or transaction channels) whose transfer of value occurs off-blockchain. If Bitcoin transactions can be signed with a new sighash type that addresses malleability, these transfers may occur between untrusted parties along the transfer route by contracts which, in the event of uncooperative or hostile participants, are enforceable via broadcast over the bitcoin blockchain in the event of uncooperative or hostile participants, through a series of decrementing timelocks.

Introduction

Commerce on the Internet has come to rely almost exclusively on financial institutions serving as trusted third parties to process electronic payments. While the system works well enough for most transactions, it still suffers from the inherent weaknesses of the trust based model. Completely non-reversible transactions are not really possible, since financial institutions cannot avoid mediating disputes. The cost of mediation increases transaction costs, limiting the minimum practical transaction size and cutting off the possibility for small casual transactions, and there is a broader cost in the loss of ability to make non-reversible payments for non-reversible services. With the possibility of reversal, the need for trust spreads. Merchants must be wary of their customers, hassling them for more information than they would otherwise need. A certain percentage of fraud is accepted as unavoidable. These costs and payment uncertainties can be avoided in person by using physical currency, but no mechanism exists to make payments over a communications channel without a trusted party.

What is needed is an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party. Transactions that are computationally impractical to reverse would protect sellers from fraud, and routine escrow mechanisms could easily be implemented to protect buyers. In this paper, we propose a solution to the double-spending problem using a peer-to-peer distributed timestamp server to generate computational proof of the chronological order of transactions. The system is secure as long as honest nodes collectively control more CPU power than any cooperating group of attacker nodes.

Interest and research in distributed consensus systems has increased markedly in recent years, with a central focus being on distributed payment networks. Such networks allow for fast, low-cost transactions which are not controlled by a centralized source. While the economic benefits and drawbacks of such a system are worthy of much research in and of themselves, this work focuses on some of the technical challenges that all distributed payment systems must face. While these problems are varied, we group them into three main categories: correctness, agreement, and utility.

By correctness, we mean that it is necessary for a distributed system to be able to discern the difference between a correct and fraudulent transaction. In traditional fiduciary settings, this is done through trust between institutions and cryptographic signatures that guarantee a transaction is indeed coming from the institution that it claims to be coming from. In distributed systems, however, there is no such trust, as the identity of any and all members in the network may not even be known. Therefore, alternative methods for

correctness must be utilized.

Agreement refers to the problem of maintaining a single global truth in the face of a decentralized accounting system. While similar to the correctness problem, the difference lies in the fact that while a malicious user of the network may be unable to create a fraudulent transaction (defying correctness), it may be able to create multiple correct transactions that are somehow unaware of each other, and thus combine to create a fraudulent act. For example, a malicious user may make two simultaneous purchases, with only enough funds in their account to cover each purchase individually, but not both together. Thus each transaction by itself is correct, but if executed simultaneously in such a way that the distributed network as a whole is unaware of both, a clear problem arises, commonly referred to as the "Double-Spend Problem." Thus the agreement problem can be summarized as the requirement that only one set of globally recognized transactions exist in the network.

Utility is a slightly more abstract problem, which we define generally as the "usefulness" of a distributed payment system, but which in practice most often simplifies to the latency of the system. A distributed system that is both correct and in agreement but which requires one year to process a transaction, for example, is obviously an inviable payment system. Additional aspects of utility may include the level of computing power required to participate in the correctness and agreement processes or the technical proficiency required of an end user to avoid being defrauded in the network.

Many of these issues have been explored long before the advent of modern distributed computer systems, via a problem known as the "Byzantine Generals Problem." In this problem, a group of generals each control a portion of an army and must coordinate an attack by sending messengers to each other. Because the generals are in unfamiliar and hostile territory, messengers may fail to reach their destination (just as nodes in a distributed network may fail, or send corrupted data instead of the intended message). An additional aspect of the problem is that some of the generals may be traitors, either individually, or conspiring together, and so messages may arrive which are intended to create a false plan that is doomed to failure for the loyal generals (just as malicious members of a distributed system may attempt to convince the system to accept fraudulent transactions, or multiple versions of the same truthful transaction that would result in a double-spend). Thus a distributed payment system must be robust both in the face of standard failures, and so-called "Byzantine" failures, which may be coordinated and originate from multiple sources in the network.

In this work, we analyze one particular implementation of a distributed payment system: the zkFUND Protocol. We focus on the algorithms utilized to achieve the above goals of correctness, agreement, and utility, and show that all are met (within necessary and predetermined tolerance thresholds, which are well-understood). In addition, we provide code that simulates the consensus process with parameterizable network size, number of malicious users, and message-sending latencies.

Satoshi Nakamoto's development of Bitcoin in 2009 has often been hailed as a radical development in money and currency, being the first example of a digital asset which simultaneously has no backing or "intrinsic value" and no centralized issuer or controller. However, another, arguably more important, part of the Bitcoin experiment is the underlying blockchain technology as a tool of distributed consensus, and attention is rapidly starting to shift to this other aspect of Bitcoin. Commonly cited alternative applications of blockchain technology include using on-blockchain digital assets to represent custom currencies and financial instruments ("colored coins"), the ownership of an underlying physical device ("smart property"), non-fungible assets such as domain names ("Namecoin"), as well as more complex applications involving having digital assets being directly controlled by a piece of code implementing arbitrary rules ("smart contracts") or even blockchain-based "decentralized autonomous organizations" (DAOs). What zkFUND intends to provide is a blockchain with a built-in fully fledged Turing-complete programming language that can be used to create "contracts" that can be used to encode arbitrary state transition functions, allowing users to create any of the

systems described above, as well as many others that we have not yet imagined, simply by writing up the logic in a few lines of code.

"Bitcoin" has been a successful implementation of the concept of p2p electronic cash. Both professionals and the general public have come to appreciate the convenient combination of public transactions and proof-of-work as a trust model. Today, the user base of electronic cash is growing at a steady pace; customers are attracted to low fees and the anonymity provided by electronic cash and merchants value its predicted and decentralized emission. Bitcoin has effectively proved that electronic cash can be as simple as paper money and as convenient as credit cards.

Unfortunately, Bitcoin suffers from several deficiencies. For example, the system's distributed nature is inflexible, preventing the implementation of new features until almost all of the network users update their clients. Some critical flaws that cannot be fixed rapidly deter Bitcoin's widespread propagation. In such inflexible models, it is more efficient to roll-out a new project rather than perpetually fix the original project.

In this paper, we study and propose solutions to the main deficiencies of Bitcoin. We believe that a system taking into account the solutions we propose will lead to a healthy competition among different electronic cash systems. We also propose our own electronic cash, "zkFUND", a name emphasizing the next breakthrough in electronic cash.

Bitcoin is the first widely used financial system for which all the necessary data to validate the system status can be cryptographically verified by anyone. However, it accomplishes this feat by storing all transactions in a public database called "the blockchain" and someone who genuinely wishes to check this state must download the whole thing and basically replay each transaction, check each one as they go. Meanwhile, most of these transactions have not affected the actual final state (they create outputs that are destroyed a transaction later).

At the time of this writing, there were nearly 150 million transactions committed in the blockchain, which must be replayed to produce a set of only 4 million unspent outputs.

It would be better if an auditor needed only to check data on the outputs themselves, but this is impossible because they are valid if and only if the output is at the end of a chain of previous outputs, each signs the next. In other words, the whole blockchain must be validated to confirm the final state.

Add to this that these transactions are cryptographically atomic, it is clear what outputs go into every transaction and what emerges. The "transaction graph" resulting reveals a lot of information and is subjected to analysis by many companies whose business model is to monitor and control the lower classes. This makes it very non-private and even dangerous for people to use.

Some solutions to this have been proposed. Greg Maxwell discovered to encrypt the amounts, so that the graph of the transaction is faceless but still allow validation that the sums are correct. Dr Maxwell also produced CoinJoin, a system for Bitcoin users to combine interactively transactions, confusing the transaction graph. Nicolas van Saberhagen has developed a system to blind the transaction entries, goes much further to cloud the transaction graph (as well as not needed the user interaction). Later, Shen Noether combined the two approaches to obtain "confidential transactions" of Maxwell AND the darkening of van Saberhagen.

These solutions are very good and would make Bitcoin very safe to use. But the problem of too much data is made even worse. Confidential transactions require multi-kilobyte proofs on every output, and van Saberhagen signatures require every output to be stored for ever, since it is not possible to tell when they are truly spent.

Dr. Maxwell's CoinJoin has the problem of needing interactivity. Dr. Yuan Horas Mouton fixed this by making transactions freely mergeable, but he needed to use pairing-based cryptography, which is potentially slower and more difficult to trust. He called this "one-way aggregate signatures" (OWAS).

OWAS had the good idea to combine the transactions in blocks. Imagine that we can combine across blocks

(perhaps with some glue data) so that when the outputs are created and destroyed, it is the same as if they never existed. Then, to validate the entire chain, users only need to know when money is entered into the system (new money in each block as in Bitcoin or Monero or peg-ins for sidechains) and final unspent outputs, the rest can be removed and forgotten.
Then we can have Confidential Transactions to hide the amounts and OWAS to blur the transaction graph, and use LESS space than Bitcoin to allow users to fully verify the blockchain. And also imagine that we must not pairing-based cryptography or new hypotheses, just regular discrete logarithms signatures like Bitcoin. Here is what I propose.
I call my creation zkFUND because it is used to prevent the blockchain from talking about all user's information.

History

The concept of decentralized digital currency, as well as alternative applications like property registries, has been around for decades. The anonymous e-cash protocols of the 1980s and the 1990s were mostly reliant on a cryptographic primitive known as Chaumian Blinding. Chaumian Blinding provided these new currencies with high degrees of privacy, but their underlying protocols largely failed to gain traction because of their reliance on a centralized intermediary. In 1998, Wei Dai's b-money became the first proposal to introduce the idea of creating money through solving computational puzzles as well as decentralized consensus, but the proposal was scant on details as to how decentralized consensus could actually be implemented. In 2005, Hal Finney introduced a concept of "reusable proofs of work", a system which uses ideas from b-money together with Adam Back's computationally difficult Hashcash puzzles to create a concept for a cryptocurrency, but once again fell short of the ideal by relying on trusted computing as a backend. In 2009, a decentralized currency was for the first time implemented in practice by Satoshi Nakamoto, combining established primitives for managing ownership through public key cryptography with a consensus algorithm for keeping track of who owns coins, known as "proof of work."
The mechanism behind proof of work was a breakthrough because it simultaneously solved two problems. First, it provided a simple and moderately effective consensus algorithm, allowing nodes in the network to collectively agree on a set of updates to the state of the Bitcoin ledger. Second, it provided a mechanism for allowing free entry into the consensus process, solving the political problem of deciding who gets to influence the consensus, while simultaneously preventing Sybil attacks. It does this by substituting a formal barrier to participation, such as the requirement to be registered as a unique entity on a particular list, with an economic barrier - the weight of a single node in the consensus voting process is directly proportional to the computing power that the node brings. Since then, an alternative approach has been proposed called proof of stake, calculating the weight of a node as being proportional to its currency holdings and not its computational resources. The discussion concerning the relative merits of the two approaches is beyond the scope of this paper but it should be noted that both approaches can be used to serve as the backbone of a cryptocurrency.

Bitcoin As A State Transition System

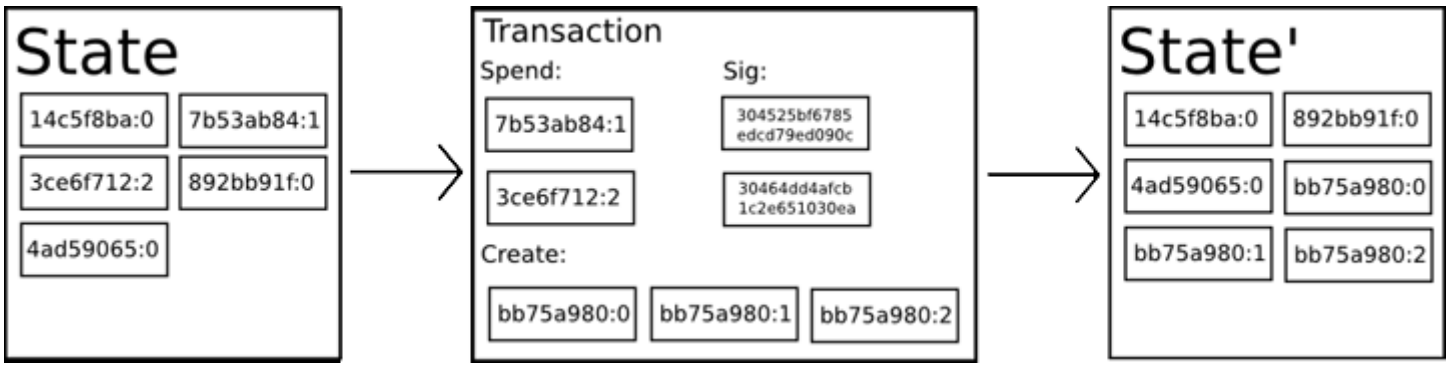

From a technical standpoint, the ledger of a cryptocurrency such as Bitcoin can be thought of as a state transition system, where there is a "state" consisting of the ownership status of all existing bitcoins and a "state transition function" that takes a state and a transaction and outputs a new state which is the result. In a standard banking system, for example, the state is a balance sheet, a transaction is a request to move $X from A to B, and the state transition function reduces the value of A's account by $X and increases the value of B's account by $X. If A's account has less than $X in the first place, the state transition function returns an error. Hence, one can formally define:

```
APPLY(S,TX) -> S' or ERROR
```

In the banking system defined above:

```
APPLY({ Alice: $50, Bob: $50 },"send $20 from Alice to Bob") = { Alice:
$30, Bob: $70 }
```

But:

```
APPLY({ Alice: $50, Bob: $50 },"send $70 from Alice to Bob") = ERROR
```

The "state" in Bitcoin is the collection of all coins (technically, "unspent transaction outputs" or UTXO) that have been minted and not yet spent, with each UTXO having a denomination and an owner (defined by a 20-byte address which is essentially a cryptographic public key). A transaction contains one or more inputs, with each input containing a reference to an existing UTXO and a cryptographic signature produced by the private key associated with the owner's address, and one or more outputs, with each output containing a new UTXO for addition to the state.

The state transition function APPLY(S,TX) -> S' can be defined roughly as follows:

1. For each input in TX: If the referenced UTXO is not in S, return an error. If the provided signature does not match the owner of the UTXO, return an error.
2. If the sum of the denominations of all input UTXO is less than the sum of the denominations of all output UTXO, return an error.
3. Return S' with all input UTXO removed and all output UTXO added.

The first half of the first step prevents transaction senders from spending coins that do not exist, the second

half of the first step prevents transaction senders from spending other people's coins, and the second step enforces conservation of value. In order to use this for payment, the protocol is as follows. Suppose Alice wants to send 11.7 BTC to Bob. First, Alice will look for a set of available UTXO that she owns that totals up to at least 11.7 BTC. Realistically, Alice will not be able to get exactly 11.7 BTC; say that the smallest she can get is 6+4+2=12. She then creates a transaction with those three inputs and two outputs. The first output will be 11.7 BTC with Bob's address as its owner, and the second output will be the remaining 0.3 BTC "change". If Alice does not claim this change by sending it to an address owned by herself, the miner will be able to claim it.

Definitions, Formalization and Previous Work

We begin by defining the components of the zkFUND Protocol. In order to prove correctness, agreement, and utility properties, we first formalize those properties into axioms. These properties, when grouped together, form the notion of consensus: the state in which nodes in the network reach correct agreement. We then highlight some previous results relating to consensus algorithms, and finally state the goals of consensus for the zkFUND Protocol within our formalization framework.

zkFUND Protocol Components

We begin our description of the zkFUND network by defining the following terms:

```
Server:
A server is any entity running the zkFUND Server software (as opposed to
the zkFUND Client software which only lets a user send and receive funds),
which participates in the consensus process.
```

```
Ledger:
The ledger is a record of the amount of currency in each userâ€™s account
and represents the "ground truth" of the network. The ledger is repeatedly
updated with transactions that successfully pass through the consensus
process.
```

```
Last-Closed Ledger:
The last-closed ledger is the most recent ledger that has been ratified by
the consensus process and thus represents the current state of the
network.
```

```
Open Ledger:
The open ledger is the current operating status of a node (each node
maintains its own open ledger). Transactions initiated by end users of a
given server are applied to the open ledger of that server, but
transactions are not considered final until they have passed through the
consensus process, at which point the open ledger becomes the last-closed
```

```
ledger.
```

```
Unique Node List (UNL):
Each server, s, main- tains a unique node list, which is a set of other servers that s queries when determining consensus. Only the votes of the other members of the UNL of s are considered when determining consensus (as opposed to every node on the network). Thus the UNL represents a subset of the network which when taken collectively, is "trusted" by s to not collude in an attempt to defraud the network. Note that this definition of "trust" does not require that each individual member of the UNL be trusted.
```

```
Proposer:
Any server can broadcast transactions to be included in the consensus process, and every server attempts to include every valid transaction when a new consensus round starts. During the consensus process, however, only proposals from servers on the UNL of a server s are considered by s.
```

Formalization

We use the term nonfaulty to refer to nodes in the net- work that behave honestly and without error. Conversely, a faulty node is one which experiences errors which may be honest (due to data corruption, implementation errors, etc.), or malicious (Byzantine errors). We reduce the notion of validating a transaction to a simple binary decision problem: each node must decide from the information it has been given on the value 0 or 1.

As in Attiya, Dolev, and Gill, 1984, we define consensus according to the following three axioms:

```
1. (C1): Every nonfaulty node makes a decision infinite time
```

```
1. (C2): All nonfaulty nodes reach the same decision value
```

```
3. (C3): 0 and 1 are both possible values for all non-faulty nodes. (This removes the trivial solution in which all nodes decide 0 or 1 regardless of the information they have been presented).
```

Existing Consensus Algorithms

There has been much research done on algorithms that achieve consensus in the face of Byzantine errors. This previous work has included extensions to cases where all participants in the network are not known ahead of time, where the messages are sent asynchronously (there is no bound on the amount of time an individual node

will take to reach a decision), and where there is a delineation between the notion of strong and weak consensus.

One pertinent result of previous work on consensus algorithms is that of Fischer, Lynch, and Patterson, 1985, which proves that in the asynchronous case, non-termination is always a possibility for a consensus algorithm, even with just one faulty process. This introduces the necessity for time-based heuristics, to ensure convergence (or at least repeated iterations of non-convergence). We shall describe these heuristics for the zkFUND Protocol later.

The strength of a consensus algorithm is usually measured in terms of the fraction of faulty processes it can tolerate. It is provable that no solution to the Byzantine Generals problem (which already assumes synchronicity, and known participants) can tolerate more than (n - 1)/3 byzantine faults, or 33% of the network acting maliciously. This solution does not, however, require verifiable authenticity of the messages delivered between nodes (digital signatures). If a guarantee on the unforgeability of messages is possible, algorithms exist with much higher fault tolerance in the synchronous case.

Several algorithms with greater complexity have been proposed for Byzantine consensus in the asynchronous case. FaB Paxos will tolerate (n - 1)/5 Byzantine failures in a network of n nodes, amounting to a tolerance of up to 20% of nodes in the network colluding maliciously. Attiya, Doyev, and Gill introduce a phase algorithm for the asynchronous case, which can tolerate (n - 1)/4 failures, or up to 25% of the network. Lastly, Alchieri et al., 2008 present BFT-CUP, which achieves Byzantine consensus in the asynchronous case even with unknown participants, with the maximal bound of a tolerance of (n - 1)/3 failures, but with additional restrictions on the connectivity of the underlying network.

Formal Consensus Goals

Our goal in this work is to show that the consensus algorithm utilized by the zkFUND Protocol will achieve consensus at each ledger-close (even if consensus is the trivial consensus of all transactions being rejected), and that the trivial consensus will only be reached with a known probability, even in the face of Byzantine failures. Our goal in this work is to show that the consensus algorithm utilized by the zkFUND Protocol will achieve consensus at each ledger-close (even if consensus is the trivial consensus of all transactions being rejected), and that the trivial consensus will only be reached with a known probability, even in the face of Byzantine failures.

Lastly we will show that the zkFUND Protocol can achieve these goals in the face of (n-1)/5 failures, which is not the strongest result in the literature, but we will also show that the zkFUND Protocol possesses several other desirable features that greatly enhance its utility.

Irregular emission

Bitcoin has a predetermined emission rate: each solved block produces a fixed amount of coins. Approximately every four years this reward is halved. The original intention was to create a limited smooth emission with exponential decay, but in fact we have a piecewise linear emission function whose breakpoints may cause problems to the Bitcoin infrastructure.

When the breakpoint occurs, miners start to receive only half of the value of their previous reward. The absolute difference between 12.5 and 6.25 BTC (projected for the year 2020) may seem tolerable. However, when examining the 50 to 25 BTC drop that took place on November 28 2012, felt inappropriate for a significant number of members of the mining community. Figure 1 shows a dramatic decrease in the network's hashrate in the end of November, exactly when the halving took place. This event could have been the perfect moment for the malevolent individual described in the proof-of-work function section to carry-out a double

spending attack.

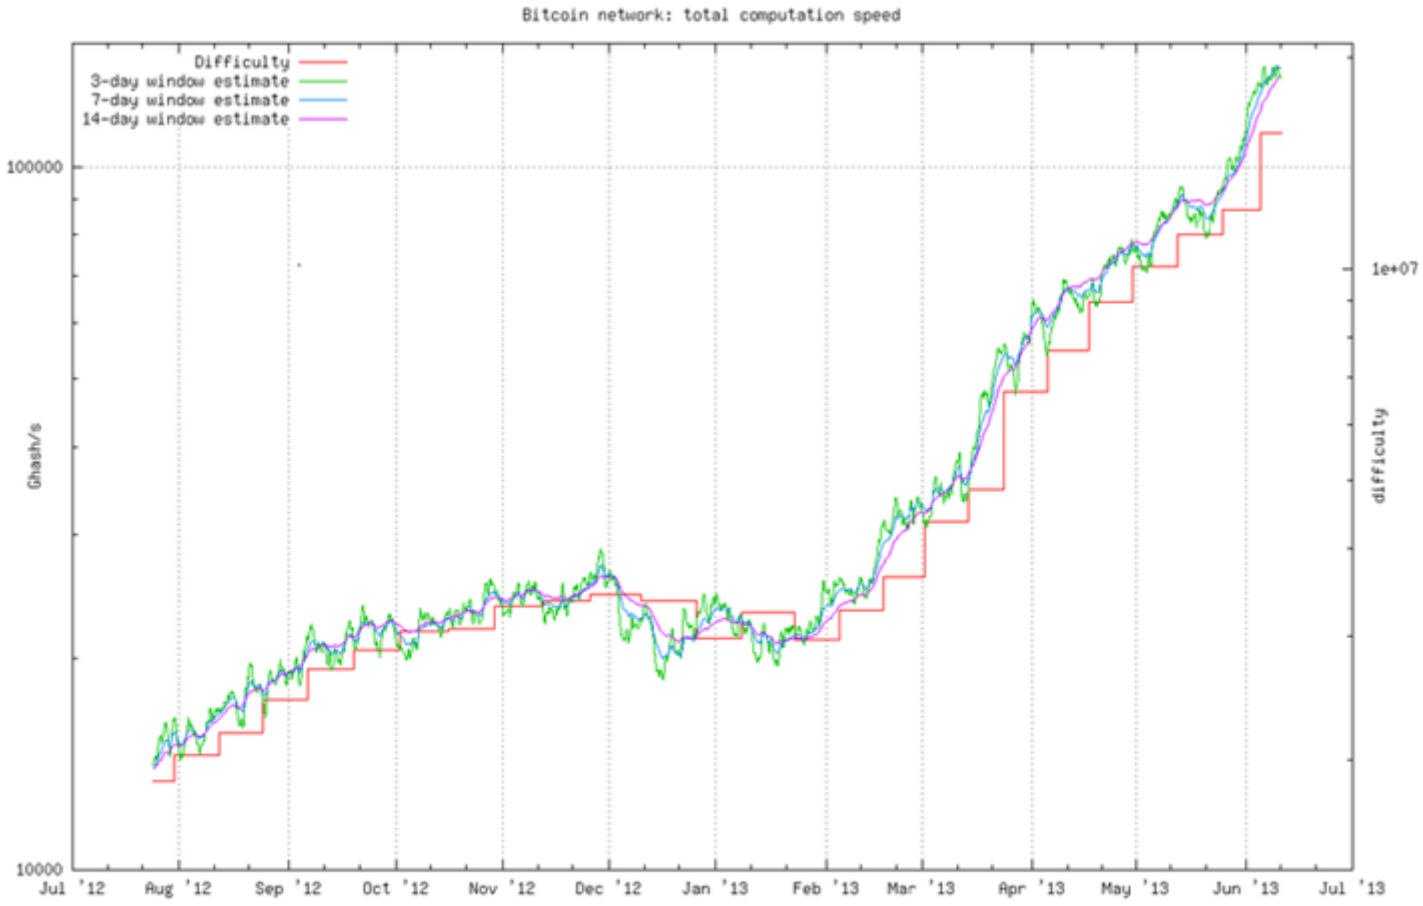

Fig. 1. Bitcoin hashrate chart
(source: http://bitcoin.sipa.be)

Hardcoded constants

Bitcoin has many hard-coded limits, where some are natural elements of the original design (e.g. block frequency, maximum amount of money supply, number of confirmations) whereas other seem to be artificial constraints. It is not so much the limits, as the inability of quickly changing them if necessary that causes the main drawbacks. Unfortunately, it is hard to predict when the constants may need to be changed and replacing them may lead to terrible consequences.

A good example of a hardcoded limit change leading to disastrous consequences is the block size limit set to 250kb1. This limit was sufficient to hold about 10000 standard transactions. In early 2013, this limit had almost been reached and an agreement was reached to increase the limit. The change was implemented in wallet version 0.8 and ended with a 24-blocks chain split and a successful double-spend attack. While the bug was not in the Bitcoin protocol, but rather in the database engine it could have been easily caught by a simple stress test if there was no artificially introduced block size limit.

Constants also act as a form of centralization point. Despite the peer-to-peer nature of Bitcoin, an overwhelming majority of nodes use the official reference client developed by a small group of people. This group makes the decision to implement changes to the protocol and most people accept these changes irrespective of their "correctness". Some decisions caused heated discussions and even calls for boycott, which indicates that the community and the developers may disagree on some important points. It therefore seems

logical to have a protocol with user-configurable and self-adjusting variables as a possible way to avoid these problems.

Bulky scripts

The scripting system in Bitcoin is a heavy and complex feature. It potentially allows one to create sophisticated transactions, but some of its features are disabled due to security concerns and some have never even been used. The script (including both senders' and receivers' parts) for the most popular transaction in Bitcoin looks like this:

```
<sig> <pubKey> OP DUP OP HASH160 <pubKeyHash> OP EQUALVERIFY OP CHECKSIG.
```

The script is 164 bytes long whereas its only purpose is to check if the receiver possess the secret key required to verify his signature.

The zkFUND Technology

Now that we have covered the limitations of the Bitcoin technology, we will concentrate on presenting the features of zkFUND.

zkFUND Consensus Algorithm

The zkFUND Protocol consensus algorithm, is applied every few seconds by all nodes, in order to maintain the correctness and agreement of the network. Once consensus is reached, the current ledger is considered "closed" and becomes the last-closed ledger. Assuming that the consensus algorithm is successful, and that there is no fork in the network, the last-closed ledger maintained by all nodes in the network will be identical.

Elliptic curve parameters

As our base signature algorithm we chose to use the fast scheme EdDSA, which is developed and implemented by D.J. Bernstein et al. Like Bitcoin's ECDSA it is based on the elliptic curve discrete logarithm problem, so our scheme could also be applied to Bitcoin in future. Common parameters are:

q: a prime number; $q = 2^{255} - 19$;

d: an element of $\mathbb{F}_q$; $d = -121665/121666$;

E: an elliptic curve equation; $-x^2 + y^2 = 1 + dx^2y^2$;

G: a base point; $G = (x, -4/5)$;

l: a prime order of the base point; $l = 2^{252} + 27742317777372353535851937790883648493$;

$\mathcal{H}_s$: a cryptographic hash function $\{0,1\}^* \to \mathbb{F}_q$;

$\mathcal{H}_p$: a deterministic hash function $E(\mathbb{F}_q) \to E(\mathbb{F}_q)$.

Terminology

Enhanced privacy requires a new terminology which should not be confused with Bitcoin entities.

private ec-key is a standard elliptic curve private key: a number $a \in [1, l-1]$;

public ec-key is a standard elliptic curve public key: a point $A = aG$;

one-time keypair is a pair of private and public ec-keys;

private user key is a pair (a, b) of two different private ec-keys;

tracking key is a pair (a, B) of private and public ec-key (where $B = bG$ and $a \neq b$);

public user key is a pair (A, B) of two public ec-keys derived from (a, b);

standard address is a representation of a public user key given into human friendly string with error correction;

truncated address is a representation of the second half (point B) of a public user key given into human friendly string with error correction.

The transaction structure remains similar to the structure in Bitcoin: every user can choose several independent incoming payments (transactions outputs), sign them with the corresponding private keys and send them to different destinations.

Contrary to Bitcoin's model, where a user possesses unique private and public key, in the proposed model a sender generates a one-time public key based on the recipient's address and some random data. In this sense, an incoming transaction for the same recipient is sent to a one-time public key (not directly to a unique address) and only the recipient can recover the corresponding private part to redeem his funds (using his unique private key). The recipient can spend the funds using a ring signature, keeping his ownership and actual spending anonymous. The details of the protocol are explained in the next subsections.

Unlinkable payments

Classic Bitcoin addresses, once being published, become unambiguous identifier for incoming payments, linking them together and tying to the recipient's pseudonyms. If someone wants to receive an "untied" transaction, he should convey his address to the sender by a private channel. If he wants to receive different transactions which cannot be proven to belong to the same owner he should generate all the different addresses and never publish them in his own pseudonym.

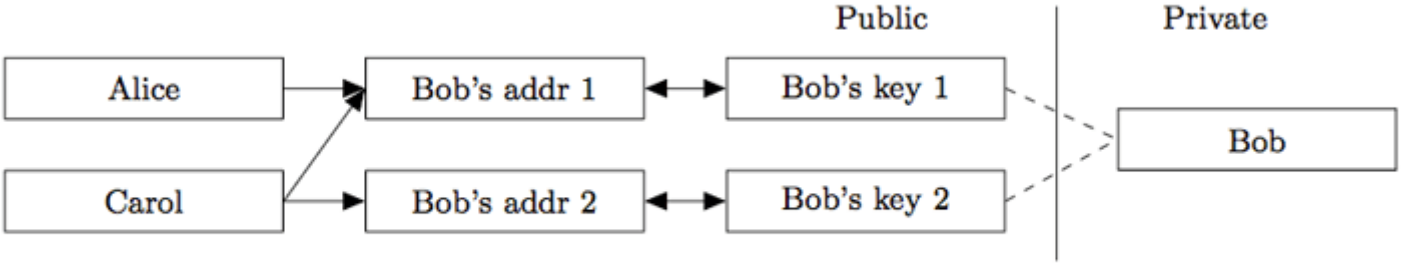

Fig. 2. Traditional Bitcoin keys/transactions model.

We propose a solution allowing a user to publish a single address and receive unconditional unlinkable payments. The destination of each zkFUND output (by default) is a public key, derived from recipient's address and sender's random data. The main advantage against Bitcoin is that every destination key is unique by default (unless the sender uses the same data for each of his transactions to the same recipient). Hence, there is no such issue as "address reuse" by design and no observer can determine if any transactions were sent to a specific address or link two addresses together.

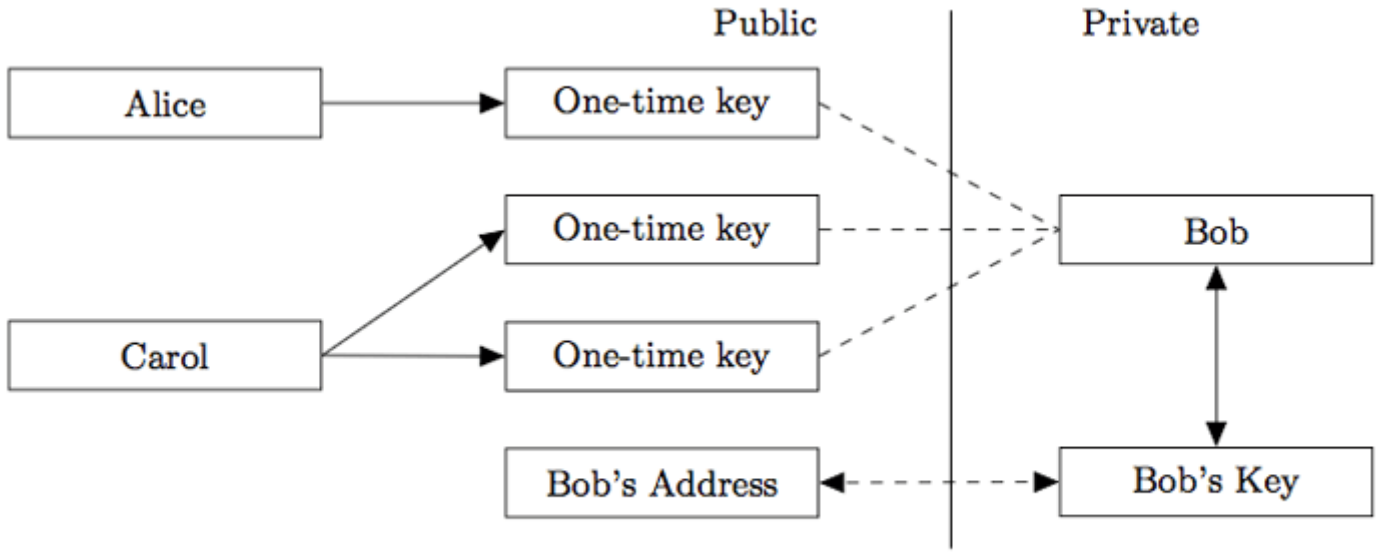

Fig. 3. CryptoNote keys/transactions model.

First, the sender performs a Diffie-Hellman exchange to get a shared secret from his data and half of the recipient's address. Then he computes a one-time destination key, using the shared secret and the second half of the address. Two different ec-keys are required from the recipient for these two steps, so a standard zkFUND address is nearly twice as large as a Bitcoin wallet address. The receiver also performs a Diffie-Hellman exchange to recover the corresponding secret key.
A standard transaction sequence goes as follows:

1. Alice wants to send a payment to Bob, who has published his standard address. She unpacks the address and gets Bob's public key (A, B).

2. Alice generates a random $r \in [1, l-1]$ and computes a one-time public key $P = \mathcal{H}_s(rA)G + B$.

3. Alice uses P as a destination key for the output and also packs value $R = rG$ (as a part of the Diffie-Hellman exchange) somewhere into the transaction. Note that she can create other outputs with unique public keys: different recipients' keys (A_i, B_i) imply different P_i even with the same r.

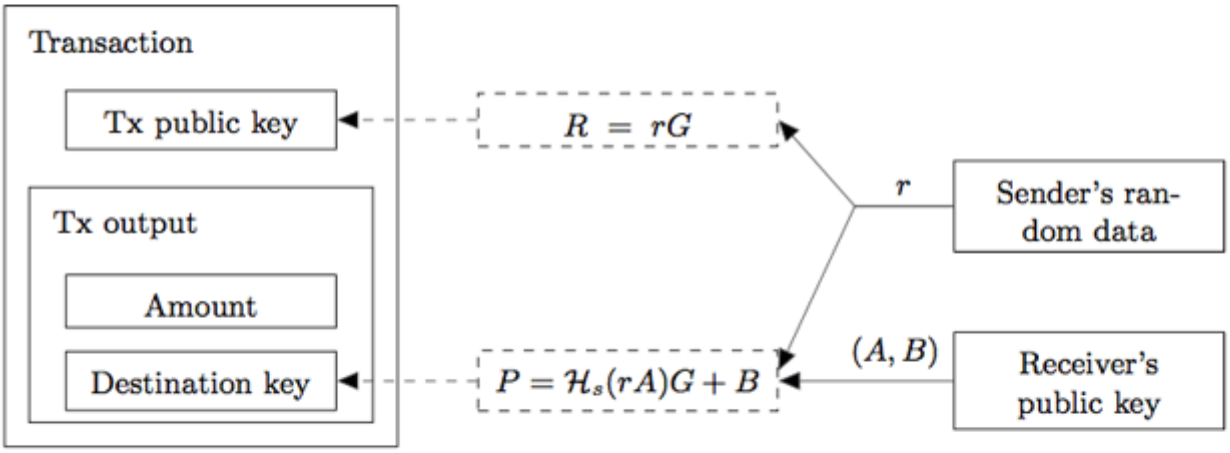

Fig. 4. Standard transaction structure.

4. Alice sends the transaction.

5. Bob checks every passing transaction with his private key (a, b), and computes $P' = \mathcal{H}_s(aR)G + B$. If Alice's transaction for with Bob as the recipient was among them, then $aR = arG = rA$ and $P' = P$.

6. Bob can recover the corresponding one-time private key: $x = \mathcal{H}_s(aR) + b$, so as $P = xG$. He can spend this output at any time by signing a transaction with x.

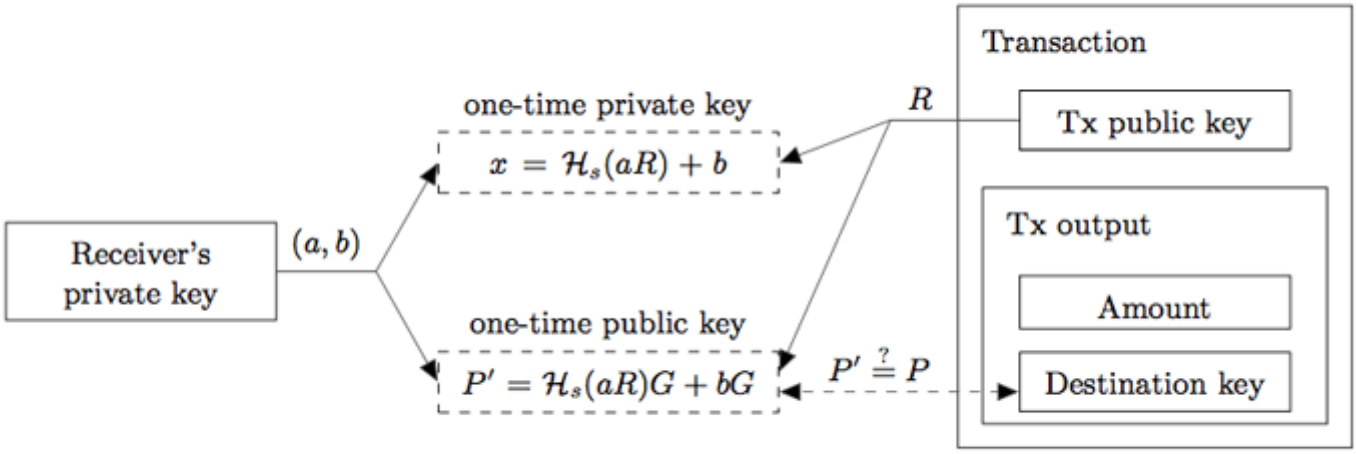

Fig. 5. Incoming transaction check.

As a result Bob gets incoming payments, associated with one-time public keys which are unlinkable for a

spectator. Some additional notes:

```
When Bob "recognizes" his transactions (see step 5) he practically uses only half of his private information: (a,B). This pair, also known as the tracking key, can be passed to a third party (Carol). Bob can delegate her the processing of new transactions. Bob doesn't need to explicitly trust Carol, because she can't recover the one-time secret key p without Bob's full private key (a, b). This approach is useful when Bob lacks bandwidth or computation power (smartphones, hardware wallets etc.).
```

```
In case Alice wants to prove she sent a transaction to Bob's address she can either disclose r or use any kind of zero-knowledge protocol to prove she knows r (for example by signing the transaction with r).
```

```
If Bob wants to have an audit compatible address where all incoming transaction are linkable, he can either publish his tracking key or use a truncated address. That address represent only one public ec-key B, and the remaining part required by the protocol is derived from it as follows: a = Hs(B) and A = Hs(B)G. In both cases every person is able to "recognize" all of Bob's incoming transaction, but, of course, none can spend the funds enclosed within them without the secret key b.
```

One-time ring signatures

A protocol based on one-time ring signatures allows users to achieve unconditional unlinkability. Unfortunately, ordinary types of cryptographic signatures permit to trace transactions to their respective senders and receivers. Our solution to this deficiency lies in using a different signature type than those currently used in electronic cash systems.

We will first provide a general description of our algorithm with no explicit reference to electronic cash.

A one-time ring signature contains four algorithms: (GEN, SIG, VER, LNK):

GEN: takes public parameters and outputs an ec-pair (P, x) and a public key I.

SIG: takes a message m, a set $\mathcal{S}'$ of public keys $\{P_i\}_{i \neq s}$, a pair (P_s, x_s) and outputs a signature σ and a set $\mathcal{S} = S' \cup \{P_s\}$.

VER: takes a message m, a set $\mathcal{S}$, a signature σ and outputs "true" or "false".

LNK: takes a set $\mathcal{I} = \{I_i\}$, a signature σ and outputs "linked" or "indep".

The idea behind the protocol is fairly simple: a user produces a signature which can be checked by a set of public keys rather than a unique public key. The identity of the signer is indistinguishable from the other users

whose public keys are in the set until the owner produces a second signature using the same keypair.

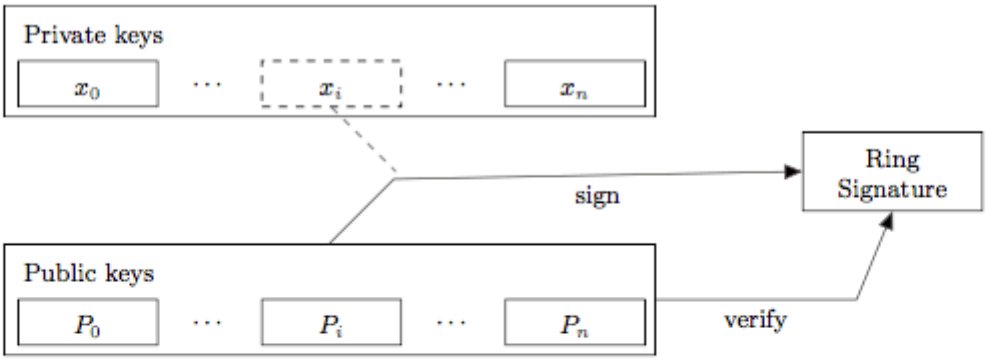

Fig. 6. Ring signature anonymity.

GEN: The signer picks a random secret key $x \in [1, l-1]$ and computes the corresponding public key $P = xG$. Additionally he computes another public key $I = x\mathcal{H}_p(P)$ which we will call the "key image".

SIG: The signer generates a one-time ring signature with a non-interactive zero-knowledge proof using the techniques from [21]. He selects a random subset $\mathcal{S}'$ of n from the other users' public keys P_i, his own keypair (x, P) and key image I. Let $0 \leq s \leq n$ be signer's secret index in $\mathcal{S}$ (so that his public key is P_s).

He picks a random $\{q_i \mid i = 0 \ldots n\}$ and $\{w_i \mid i = 0 \ldots n, i \neq s\}$ from $(1 \ldots l)$ and applies the following *transformations*:

$$L_i = \begin{cases} q_i G, & \text{if } i = s \\ q_i G + w_i P_i, & \text{if } i \neq s \end{cases}$$

$$R_i = \begin{cases} q_i \mathcal{H}_p(P_i), & \text{if } i = s \\ q_i \mathcal{H}_p(P_i) + w_i I, & \text{if } i \neq s \end{cases}$$

The next step is getting the non-interactive *challenge*:

$$c = \mathcal{H}_s(m, L_1, \ldots, L_n, R_1, \ldots, R_n)$$

Finally the signer computes the *response*:

$$c_i = \begin{cases} w_i, & \text{if } i \neq s \\ c - \sum_{i=0}^{n} c_i \mod l, & \text{if } i = s \end{cases}$$

$$r_i = \begin{cases} q_i, & \text{if } i \neq s \\ q_s - c_s x \mod l, & \text{if } i = s \end{cases}$$

The resulting signature is $\sigma = (I, c_1, \ldots, c_n, r_1, \ldots, r_n)$.

VER: The verifier checks the signature by applying the inverse transformations:

$$\begin{cases} L'_i = r_i G + c_i P_i \\ R'_i = r_i \mathcal{H}_p(P_i) + c_i I \end{cases}$$

Finally, the verifier checks if $\sum_{i=0}^{n} c_i \stackrel{?}{=} \mathcal{H}_s(m, L'_0, \ldots, L'_n, R'_0, \ldots, R'_n) \mod l$

If this equality is correct, the verifier runs the algorithm **LNK**. Otherwise the verifier rejects the signature.

LNK: The verifier checks if I has been used in past signatures (these values are stored in the set $\mathcal{I}$). Multiple uses imply that two signatures were produced under the same secret key.

The meaning of the protocol: by applying L-transformations the signer proves that he knows such x that at least one $P_i = xG$. To make this proof non-repeatable we introduce the key image as $I = x\mathcal{H}_p(P)$. The signer uses the same coefficients (r_i, c_i) to prove almost the same statement: he knows such x that at least one $\mathcal{H}_p(P_i) = I \cdot x^{-1}$.

If the mapping $x \to I$ is an injection:

1. Nobody can recover the public key from the key image and identify the signer;
2. The signer cannot make two signatures with different I's and the same x.

Standard zkFUND transaction

By combining both methods (unlinkable public keys and untraceable ring signature) Bob achieves new level of privacy in comparison with the original Bitcoin scheme. It requires him to store only one private key (a, b) and publish (A, B) to start receiving and sending anonymous transactions.

While validating each transaction Bob additionally performs only two elliptic curve multiplications and one addition per output to check if a transaction belongs to him. For his every output Bob recovers a one-time keypair (pi,Pi) and stores it in his wallet. Any inputs can be circumstantially proved to have the same owner only if they appear in a single transaction. In fact this relationship is much harder to establish due to the one-time ring signature.

With a ring signature Bob can effectively hide every input among somebody else's; all possible spenders will be equiprobable, even the previous owner (Alice) has no more information than any observer.

When signing his transaction Bob specifies n foreign outputs with the same amount as his output, mixing all of them without the participation of other users. Bob himself (as well as anybody else) does not know if any of these payments have been spent: an output can be used in thousands of signatures as an ambiguity factor and never as a target of hiding. The double spend check occurs in the LNK phase when checking against the used key images set.

Bob can choose the ambiguity degree on his own: n = 1 means that the probability he has spent the output is 50% probability, n = 99 gives 1%. The size of the resulting signature increases linearly as O(n + 1), so the improved anonymity costs to Bob extra transaction fees. He also can set n = 0 and make his ring signature to consist of only one element, however this will instantly reveal him as a spender.

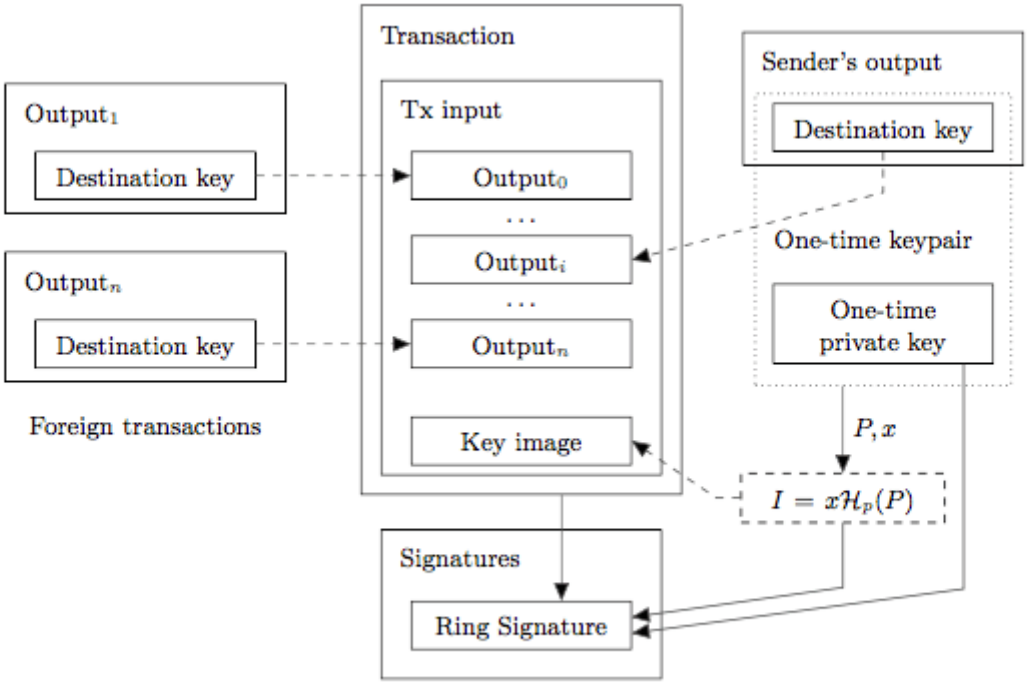

Fig. 7. Ring signature generation in a standard transaction.

Egalitarian Proof-of-work

In this section we propose and ground the new proof-of-work algorithm. Our primary goal is to close the gap between CPU (majority) and GPU/FPGA/ASIC (minority) miners. It is appropriate that some users can have a certain advantage over others, but their investments should grow at least linearly with the power. More generally, producing special-purpose devices has to be as less profitable as possible.

Definition

The RPCA proceeds in rounds. In each round:

```
Initially, each server takes all valid transactions it has seen prior to
the beginning of the consensus round that have not already been applied
(these may include new transactions initiated by endusers of the server,
transactions held over from a previous consensus process, etc.), and makes
them public in the form of a list known as the "candidate set".
```

```
Each server then amalgamates the candidate sets of all servers on its UNL,
and votes on the veracity of all transactions.
```

```
Transactions that receive more than a minimum percentage of "yes" votes
are passed on to the next round, if there is one, while transactions that
do not receive enough votes will either be discarded, or included in the
candidate set for the beginning of the consensus process on the next
ledger.
```

```
The final round of consensus requires a minimum percentage of 80% of a
server's UNL agreeing on a transaction. All transactions that meet this
requirement are applied to the ledger, and that ledger is closed, becoming
the new last-closed ledger.
```

Transactions

We define an electronic coin as a chain of digital signatures. Each owner transfers the coin to the next by digitally signing a hash of the previous transaction and the public key of the next owner and adding these to the end of the coin. A payee can verify the signatures to verify the chain of ownership.

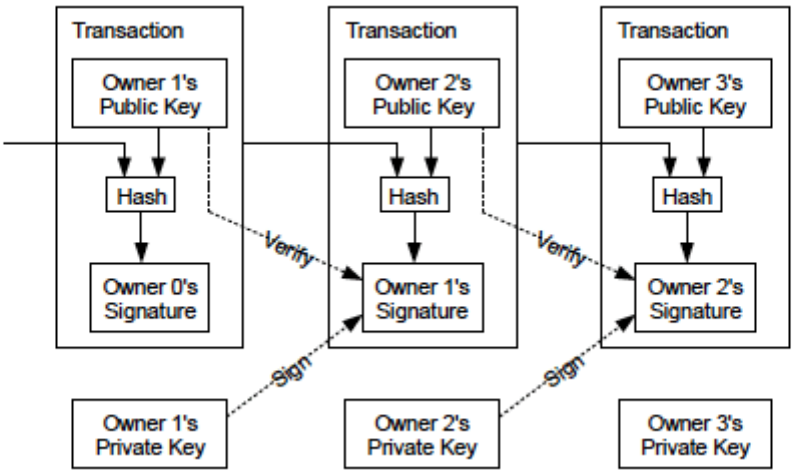

The problem of course is the payee can't verify that one of the owners did not double-spend the coin. A common solution is to introduce a trusted central authority, or mint, that checks every transaction for double spending. After each transaction, the coin must be returned to the mint to issue a new coin, and only coins issued directly from the mint are trusted not to be double-spent. The problem with this solution is that the fate of the entire money system depends on the company running the mint, with every transaction having to go through them, just like a bank.

We need a way for the payee to know that the previous owners did not sign any earlier transactions. For our purposes, the earliest transaction is the one that counts, so we don't care about later attempts to double-spend. The only way to confirm the absence of a transaction is to be aware of all transactions. In the mint based model, the mint was aware of all transactions and decided which arrived first. To accomplish this without a trusted party, transactions must be publicly announced, and we need a system for participants to agree on a single history of the order in which they were received. The payee needs proof that at the time of each transaction, the majority of nodes agreed it was the first received.

Confidential Transactions and OWAS

The first thing we need to do is remove Bitcoin Script. This is sad, but it is too powerful so it is impossible to merge transactions using general scripts. We will demonstrate that confidential transactions of Dr. Maxwell are enough (after some small modification) to authorize spending of outputs and also allows to make combined transactions without interaction. This is in fact identical to OWAS, and allows relaying nodes take some transaction fee or the recipient to change the transaction fees. These additional things Bitcoin can not do, we get for free.

We start by reminding the reader how confidential transactions work. First, the amounts are coded by the following equation:

```
C = r*G + v*H
```

where C is a Pedersen commitment, G and H are fixed nothing-up-my-sleeve elliptic curve group generators, v is the amount, and r is a secret random blinding key.
Attached to this output is a rangeproof which proves that v is in [0, 2^64], so that user cannot exploit the blinding to produce overflow attacks, etc.
To validate a transaction, the verifier will add commitments for all outputs, plus f*H (f here is the transaction fee which is given explicitly) and subtracts all input commitments. The result must be 0, which proves that no amount was created or destroyed overall.
We note that to create such a transaction, the user must know the sum of all the values of r for commitments entries. Therefore, the r-values (and their sums) act as secret keys. If we can make the r output values known only to the recipient, then we have an authentication system! Unfortunately, if we keep the rule that commits all add to 0, this is impossible, because the sender knows the sum of all _his_ r values, and therefore knows the receipient's r values sum to the negative of that. So instead, we allow the transaction to sum to a nonzero value k*G, and require a signature of an empty string with this as key, to prove its amount component is zero.
We let transactions have as many k*G values as they want, each with a signature, and sum them during verification.

To create transactions sender and recipient do following ritual:

```
1. Sender and recipient agree on amount to be sent. Call this b.

2. Sender creates transaction with all inputs and change output(s), and
gives recipient the total blinding factor (r-value of change minus
r-values of inputs) along with this transaction. So the commitments sum to
r*G - b*H.

3. Recipient chooses random r-values for his outputs, and values that sum
to b minus fee, and adds these to transaction (including range proof).
Now the commitments sum to k*G - fee*H for some k that only recipient
knows.

4. Recipient attaches signature with k to the transaction, and the
explicit fee. It has done.
```

Now, creating transactions in this manner supports OWAS already. To show this, suppose we have two transactions that have a surplus k1*G and k2*G, and the attached signatures with these. Then you can combine the lists of inputs and outputs of the two transactions, with both k1*G and k2*G to the mix, and voilÃ¡! is again a valid transaction. From the combination, it is impossible to say which outputs or inputs are from which original transaction.

Because of this, we change our block format from Bitcoin to this information:

```
1. Explicit amounts for new money (block subsidy or sidechain peg-ins)
with whatever else data this needs. For a sidechain peg-in maybe it
references a Bitcoin transaction that commits to a specific excess k*G
value?

2. Inputs of all transactions

3. Outputs of all transactions

4. Excess k*G values for all transactions
```

Each of these are grouped together because it do not matter what the transaction boundaries are originally. In addition, Lists 2 3 and 4 should be required to be coded in alphabetical order, since it is quick to check and prevents the block creator of leaking any information about the original transactions.
Note that the outputs are now identified by their hash, and not by their position in a transaction that could easily change. Therefore, it should be banned to have two unspent outputs are equal at the same time, to avoid confusion.

Traceability of Transactions

Privacy and anonymity are the most important aspects of electronic cash. Peer-to-peer payments seek to be concealed from third party's view, a distinct difference when compared with traditional banking. In particular, T. Okamoto and K. Ohta described six criteria of ideal electronic cash, which included "privacy: relationship between the user and his purchases must be untraceable by anyone." From their description, we derived two properties which a fully anonymous electronic cash model must satisfy in order to comply with the requirements outlined by Okamoto and Ohta:

```
Untraceability: for each incoming transaction all possible senders are
equiprobable.
```

```
Unlinkability: for any two outgoing transactions it is impossible to prove
they were sent to the same person.
```

Unfortunately, Bitcoin does not satisfy the untraceability requirement. Since all the transactions that take place between the network's participants are public, any transaction can be unambiguously traced to a unique origin and final recipient. Even if two participants exchange funds in an indirect way, a properly engineered path-finding method will reveal the origin and final recipient.
It is also suspected that Bitcoin does not satisfy the second property. Some researchers stated that a careful blockchain analysis may reveal a connection between the users of the Bitcoin network and their transactions. Although a number of methods are disputed, it is suspected that a lot of hidden personal information can be extracted from the public database.
Bitcoin's failure to satisfy the two properties outlined above leads us to conclude that it is not an anonymous

but a pseudo-anonymous electronic cash system. Users were quick to develop solutions to circumvent this shortcoming. Two direct solutions were "laundering services" and the development of distributed methods. Both solutions are based on the idea of mixing several public transactions and sending them through some intermediary address; which in turn suffers the drawback of requiring a trusted third party.
Recently, a more creative scheme was proposed by I. Miers et al.: "Zerocoin". Zerocoin utilizes a cryptographic one-way accumulators and zero-knoweldge proofs which permit users to "convert" bitcoins to zerocoins and spend them using anonymous proof of ownership instead of explicit public-key based digital signatures. However, such knowledge proofs have a constant but inconvenient size - about 30kb (based on today's Bitcoin limits), which makes the proposal impractical. Authors admit that the protocol is unlikely to ever be accepted by the majority of Bitcoin users.

Untraceable Transactions

In this section we propose a scheme of fully anonymous transactions satisfying both untraceability and unlinkability conditions. An important feature of our solution is its autonomy: the sender is not required to cooperate with other users or a trusted third party to make his transactions; hence each participant produces a cover traffic independently.

Merging Transactions Across Blocks

Now, we have used Dr. Maxwell's Confidential Transactions to create a noninteractive version of Dr. Maxwell's CoinJoin, but we have not seen the last of marvelous Dr. Maxwell! We need another idea, transaction cut-through, he described in. Again, we create a noninteractive version of this, and to show how it is used with several blocks.
We can imagine now each block as one large transaction. To validate it, we add all the output commitments together, then subtracts all input commitments, k*G values, and all explicit input amounts times H. We find that we could combine transactions from two blocks, as we combined transactions to form a single block, and the result is again a valid transaction. Except now, some output commitments have an input commitment exactly equal to it, where the first block's output was spent in the second block. We could remove both commitments and still have a valid transaction. In fact, there is not even need to check the rangeproof of the deleted output.
The extension of this idea all the way from the genesis block to the latest block, we see that EVERY nonexplicit input is deleted along with its referenced output. What remains are only the unspent outputs, explicit input amounts and every k*G value. And this whole mess can be validated as if it were one transaction: add all unspent commitments output, subtract the values k*G, validate explicit input amounts (if there is anything to validate) then subtract them times H. If the sum is 0, the entire chain is good.

What is this mean? When a user starts up and downloads the chain he needs the following
data from each block:

```
1. Explicit amounts for new money (block subsidy or sidechain peg-ins)
with whatever else data this needs.

2. Unspent outputs of all transactions, along with a merkle proof that
each output appeared in the original block.
```

3. Excess k*G values for all transactions.

Bitcoin today there are about 423000 blocks, totaling 80GB or so of data on the hard drive to validate everything. These data are about 150 million transactions and 5 million unspent nonconfidential outputs. Estimate how much space the number of transactions take on a zkFUND chain. Each unspent output is around 3Kb for rangeproof and Merkle proof. Each transaction also adds about 100 bytes: a k*G value and a signature. The block headers and explicit amounts are negligible. Add this together and get 30Gb -- with a confidential transaction and obscured transaction graph!

Timestamp Server

The solution we propose begins with a timestamp server. A timestamp server works by taking a hash of a block of items to be timestamped and widely publishing the hash, such as in a newspaper or Usenet post. The timestamp proves that the data must have existed at the time, obviously, in order to get into the hash. Each timestamp includes the previous timestamp in its hash, forming a chain, with each additional timestamp reinforcing the ones before it.

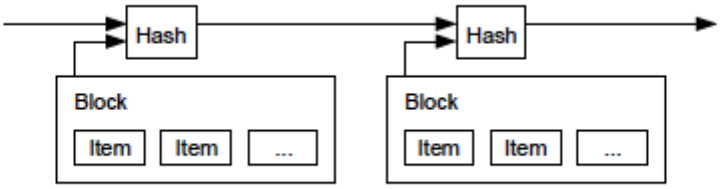

Proof-of-Work

To implement a distributed timestamp server on a peer-to-peer basis, we will need to use a proof- of-work system similar to Adam Back's Hashcash, rather than newspaper or Usenet posts. The proof-of-work involves scanning for a value that when hashed, such as with SHA-256, the hash begins with a number of zero bits. The average work required is exponential in the number of zero bits required and can be verified by executing a single hash.

For our timestamp network, we implement the proof-of-work by incrementing a nonce in the block until a value is found that gives the block's hash the required zero bits. Once the CPU effort has been expended to make it satisfy the proof-of-work, the block cannot be changed without redoing the work. As later blocks are chained after it, the work to change the block would include redoing all the blocks after it.

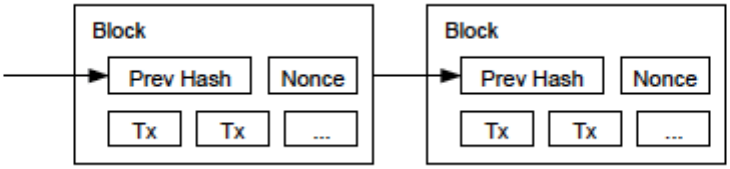

The proof-of-work also solves the problem of determining representation in majority decision making. If the majority were based on one-IP-address-one-vote, it could be subverted by anyone able to allocate many IPs. Proof-of-work is essentially one-CPU-one-vote. The majority decision is represented by the longest chain, which has the greatest proof-of-work effort invested in it. If a majority of CPU power is controlled by honest

nodes, the honest chain will grow the fastest and outpace any competing chains. To modify a past block, an attacker would have to redo the proof-of-work of the block and all blocks after it and then catch up with and surpass the work of the honest nodes. We will show later that the probability of a slower attacker catching up diminishes exponentially as subsequent blocks are added.

To compensate for increasing hardware speed and varying interest in running nodes over time, the proof-of-work difficulty is determined by a moving average targeting an average number of blocks per hour. If they're generated too fast, the difficulty increases.

Bitcoin creator Satoshi Nakamoto described the majority decision making algorithm as "one-CPU-one-vote" and used a CPU-bound pricing function (double SHA-256) for his proof-of-work scheme. Since users vote for the single history of transactions order, the reasonableness and consistency of this process are critical conditions for the whole system.

The security of this model suffers from two drawbacks. First, it requires 51% of the network's mining power to be under the control of honest users. Secondly, the system's progress (bug fixes, security fixes, etc...) require the overwhelming majority of users to support and agree to the changes (this occurs when the users update their wallet software).Finally this same voting mechanism is also used for collective polls about implementation of some features.

This permits us to conjecture the properties that must be satisfied by the proof-of-work pricing function. Such function must not enable a network participant to have a significant advantage over another participant; it requires a parity between common hardware and high cost of custom devices. From recent examples, we can see that the SHA-256 function used in the Bitcoin architecture does not posses this property as mining becomes more efficient on GPUs and ASIC devices when compared to high-end CPUs.

Therefore, Bitcoin creates favourable conditions for a large gap between the voting power of participants as it violates the "one-CPU-one-vote" principle since GPU and ASIC owners posses a much larger voting power when compared with CPU owners. It is a classical example of the Pareto principle where 20% of a system's participants control more than 80% of the votes.

One could argue that such inequality is not relevant to the network's security since it is not the small number of participants controlling the majority of the votes but the honesty of these participants that matters. However, such argument is somewhat flawed since it is rather the possibility of cheap specialized hardware appearing rather than the participants' honesty which poses a threat. To demonstrate this, let us take the following example. Suppose a malevolent individual gains significant mining power by creating his own mining farm through the cheap hardware described previously. Suppose that the global hashrate decreases significantly, even for a moment, he can now use his mining power to fork the chain and double-spend. As we shall see later in this article, it is not unlikely for the previously described event to take place.

Network

The steps to run the network are as follows:

1) New transactions are broadcast to all nodes.
2) Each node collects new transactions into a block.
3) Each node works on finding a difficult proof-of-work for its block.
4) When a node finds a proof-of-work, it broadcasts the block to all nodes.
5) Nodes accept the block only if all transactions in it are valid and not already spent.
6) Nodes express their acceptance of the block by working on creating the next block in the chain, using the hash of the accepted block as the previous hash.

Nodes always consider the longest chain to be the correct one and will keep working on extending it. If two nodes broadcast different versions of the next block simultaneously, some nodes may receive one or the other first. In that case, they work on the first one they received, but save the other branch in case it becomes longer. The tie will be broken when the next proof-of-work is found and one branch becomes longer; the nodes that were working on the other branch will then switch to the longer one.

New transaction broadcasts do not necessarily need to reach all nodes. As long as they reach many nodes, they will get into a block before long. Block broadcasts are also tolerant of dropped messages. If a node does not receive a block, it will request it when it receives the next block and realizes it missed one.

Incentive

By convention, the first transaction in a block is a special transaction that starts a new coin owned by the creator of the block. This adds an incentive for nodes to support the network, and provides a way to initially distribute coins into circulation, since there is no central authority to issue them. The steady addition of a constant of amount of new coins is analogous to gold miners expending resources to add gold to circulation. In our case, it is CPU time and electricity that is expended.

resources to add gold to circulation. In our case, it is CPU time and electricity that is expended. The incentive can also be funded with transaction fees. If the output value of a transaction is less than its input value, the difference is a transaction fee that is added to the incentive value of the block containing the transaction. Once a predetermined number of coins have entered circulation, the incentive can transition entirely to transaction fees and be completely inflation free.

The incentive may help encourage nodes to stay honest. If a greedy attacker is able to assemble more CPU power than all the honest nodes, he would have to choose between using it to defraud people by stealing back his payments, or using it to generate new coins. He ought to find it more profitable to play by the rules, such rules that favour him with more new coins than everyone else combined, than to undermine the system and the validity of his own wealth.

Reclaiming Disk Space

Once the latest transaction in a coin is buried under enough blocks, the spent transactions before it can be discarded to save disk space. To facilitate this without breaking the block's hash, transactions are hashed in a Merkle Tree, with only the root included in the block's hash. Old blocks can then be compacted by stubbing off branches of the tree. The interior hashes do not need to be stored.

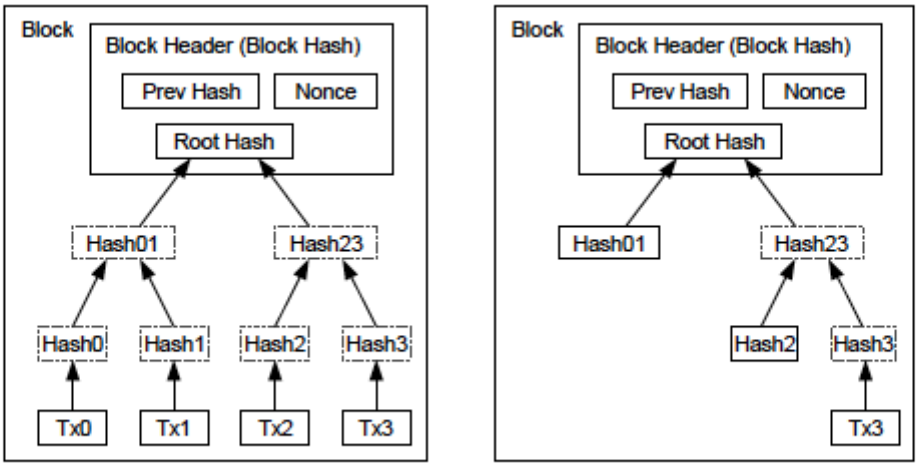

A block header with no transactions would be about 80 bytes. If we suppose blocks are generated every 10 minutes, 80 bytes * 6 * 24 * 365 = 4.2MB per year. With computer systems typically selling with 2GB of RAM as of 2008, and Moore's Law predicting current growth of 1.2GB per year, storage should not be a problem even if the block headers must be kept in memory.

Simplified Payment Verification

It is possible to verify payments without running a full network node. A user only needs to keep a copy of the block headers of the longest proof-of-work chain, which he can get by querying network nodes until he's convinced he has the longest chain, and obtain the Merkle branch linking the transaction to the block it's timestamped in. He can't check the transaction for himself, but by linking it to a place in the chain, he can see that a network node has accepted it, and blocks added after it further confirm the network has accepted it.

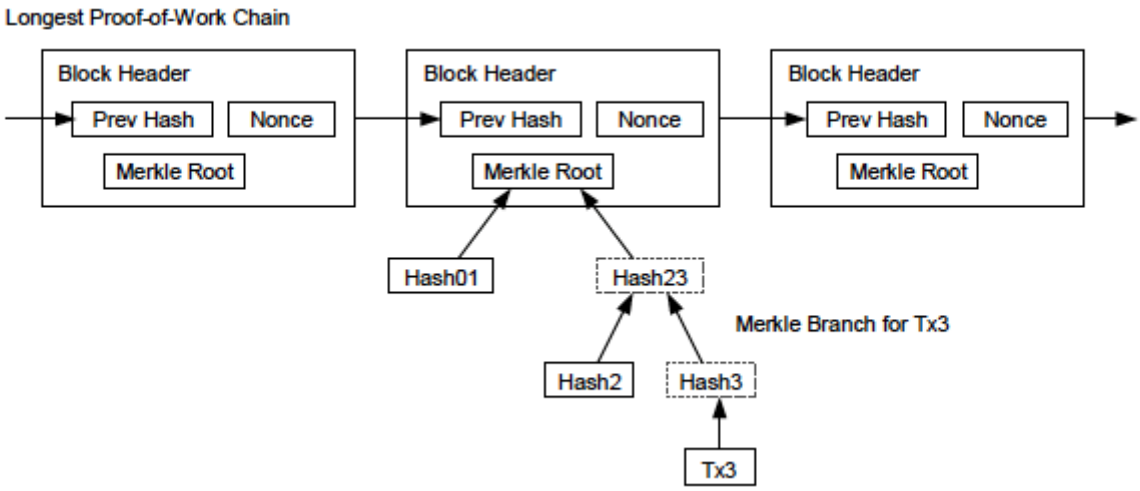

As such, the verification is reliable as long as honest nodes control the network, but is more vulnerable if the network is overpowered by an attacker. While network nodes can verify transactions for themselves, the simplified method can be fooled by an attacker's fabricated transactions for as long as the attacker can continue to overpower the network. One strategy to protect against this would be to accept alerts from network nodes when they detect an invalid block, prompting the user's software to download the full block and alerted transactions to confirm the inconsistency. Businesses that receive frequent payments will probably still want to run their own nodes for more independent security and quicker verification.

Combining and Splitting Value

Although it would be possible to handle coins individually, it would be unwieldy to make a separate transaction for every cent in a transfer. To allow value to be split and combined, transactions contain multiple inputs and outputs. Normally there will be either a single input from a larger previous transaction or multiple inputs combining smaller amounts, and at most two outputs: one for the payment, and one returning the change, if any, back to the sender.

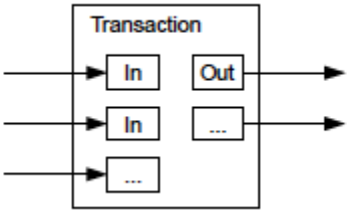

It should be noted that fan-out, where a transaction depends on several transactions, and those transactions depend on many more, is not a problem here. There is never the need to extract a complete standalone copy of a transaction's history.

Privacy

The traditional banking model achieves a level of privacy by limiting access to information to the parties involved and the trusted third party. The necessity to announce all transactions publicly precludes this method, but privacy can still be maintained by breaking the flow of information in another place: by keeping public keys anonymous. The public can see that someone is sending an amount to someone else, but without information linking the transaction to anyone. This is similar to the level of information released by stock exchanges, where the time and size of individual trades, the "tape", is made public, but without telling who the parties were.

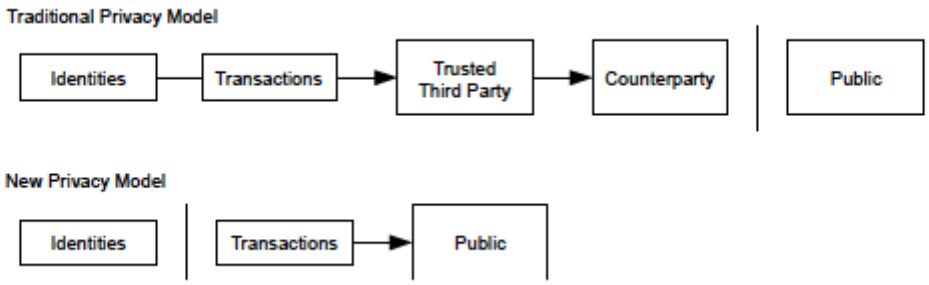

As an additional firewall, a new key pair should be used for each transaction to keep them from being linked to a common owner. Some linking is still unavoidable with multi-input transactions, which necessarily reveal that their inputs were owned by the same owner. The risk is that if the owner of a key is revealed, linking could reveal other transactions that belonged to the same owner.

Mining

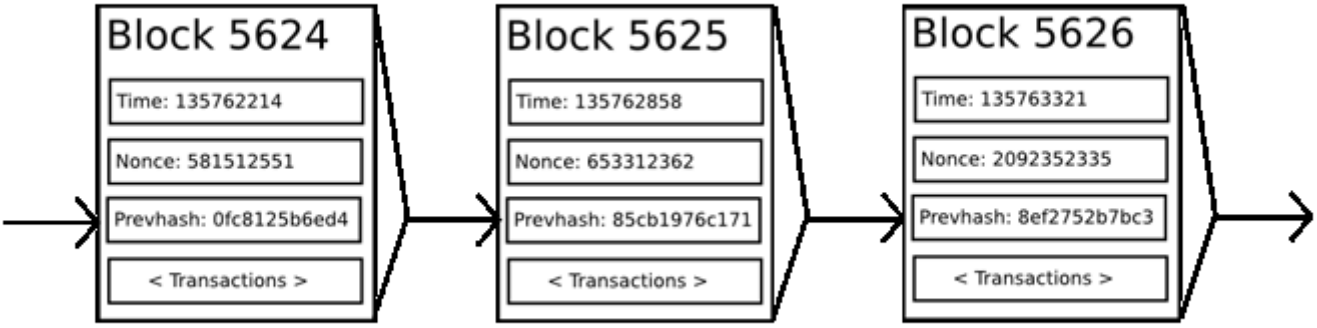

If we had access to a trustworthy centralized service, this system would be trivial to implement; it could be coded exactly as described, using a centralized server's hard drive to keep track of the state. However, with Bitcoin we are trying to build a decentralized currency system, so we will need to combine the state transition system with a consensus system in order to ensure that everyone agrees on the order of transactions. Bitcoin's decentralized consensus process requires nodes in the network to continuously attempt to produce packages of transactions called "blocks". The network is intended to create one block approximately every ten minutes, with each block containing a timestamp, a nonce, a reference to (i.e., hash of) the previous block and a list of all of the transactions that have taken place since the previous block. Over time, this creates a persistent, ever-growing, "blockchain" that continually updates to represent the latest state of the Bitcoin ledger.

The algorithm for checking if a block is valid, expressed in this paradigm, is as follows:
1. Check if the previous block referenced by the block exists and is valid.
2. Check that the timestamp of the block is greater than that of the previous block and less than 2 hours into the future
3. Check that the proof of work on the block is valid.
4. Let S[0] be the state at the end of the previous block.
5. Suppose TX is the block's transaction list with n transactions. For all i in 0...n-1, set S[i+1] = APPLY(S[i],TX[i]) If any application returns an error, exit and return false.
6. Return true, and register S[n] as the state at the end of this block.

Essentially, each transaction in the block must provide a valid state transition from what was the canonical state before the transaction was executed to some new state. Note that the state is not encoded in the block in any way; it is purely an abstraction to be remembered by the validating node and can only be (securely) computed for any block by starting from the genesis state and sequentially applying every transaction in every block. Additionally, note that the order in which the miner includes transactions into the block matters; if there are two transactions A and B in a block such that B spends a UTXO created by A, then the block will be valid if A comes before B but not otherwise.
The one validity condition present in the above list that is not found in other systems is the requirement for "proof of work". The precise condition is that the double-SHA256 hash of every block, treated as a 256-bit number, must be less than a dynamically adjusted target, which as of the time of this writing is approximately 2^187. The purpose of this is to make block creation computationally "hard", thereby preventing Sybil attackers from remaking the entire blockchain in their favor. Because SHA256 is designed to be a completely unpredictable pseudorandom function, the only way to create a valid block is simply trial and error, repeatedly

incrementing the nonce and seeing if the new hash matches.
At the current target of ~2^187, the network must make an average of ~2^69 tries before a valid block is found; in general, the target is recalibrated by the network every 2016 blocks so that on average a new block is produced by some node in the network every ten minutes. In order to compensate miners for this computational work, the miner of every block is entitled to include a transaction giving themselves 25 BTC out of nowhere. Additionally, if any transaction has a higher total denomination in its inputs than in its outputs, the difference also goes to the miner as a "transaction fee". Incidentally, this is also the only mechanism by which BTC are issued; the genesis state contained no coins at all.
In order to better understand the purpose of mining, let us examine what happens in the event of a malicious attacker. Since Bitcoin's underlying cryptography is known to be secure, the attacker will target the one part of the Bitcoin system that is not protected by cryptography directly: the order of transactions. The attacker's strategy is simple:

1. Send 100 BTC to a merchant in exchange for some product (preferably a rapid-delivery digital good)
2. Wait for the delivery of the product
3. Produce another transaction sending the same 100 BTC to himself

Try to convince the network that his transaction to himself was the one that came first.

Once step (1) has taken place, after a few minutes some miner will include the transaction in a block, say block number 270000. After about one hour, five more blocks will have been added to the chain after that block, with each of those blocks indirectly pointing to the transaction and thus "confirming" it. At this point, the merchant will accept the payment as finalized and deliver the product; since we are assuming this is a digital good, delivery is instant. Now, the attacker creates another transaction sending the 100 BTC to himself. If the attacker simply releases it into the wild, the transaction will not be processed; miners will attempt to run APPLY(S,TX) and notice that TX consumes a UTXO which is no longer in the state. So instead, the attacker creates a "fork" of the blockchain, starting by mining another version of block 270000 pointing to the same block 269999 as a parent but with the new transaction in place of the old one. Because the block data is different, this requires redoing the proof of work. Furthermore, the attacker's new version of block 270000 has a different hash, so the original blocks 270001 to 270005 do not "point" to it; thus, the original chain and the attacker's new chain are completely separate. The rule is that in a fork the longest blockchain is taken to be the truth, and so legitimate miners will work on the 270005 chain while the attacker alone is working on the 270000 chain. In order for the attacker to make his blockchain the longest, he would need to have more computational power than the rest of the network combined in order to catch up (hence, "51% attack").

Merkle Trees

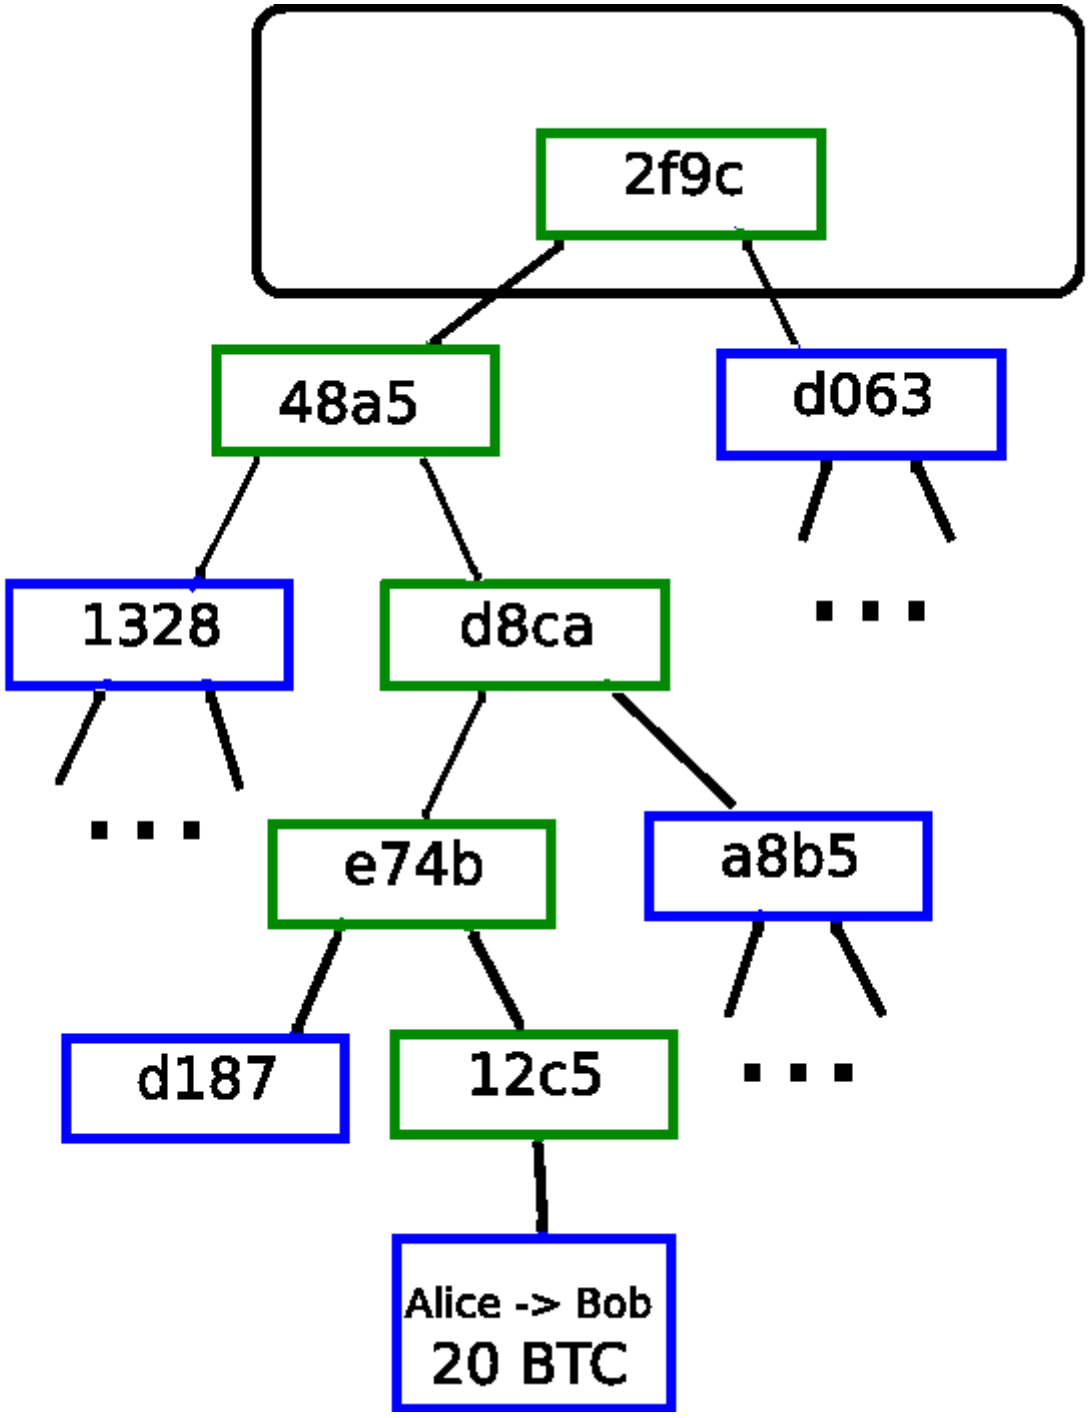


```
01: it suffices to present only a small number of nodes in a Merkle tree
to give a proof of the validity of a branch.
```

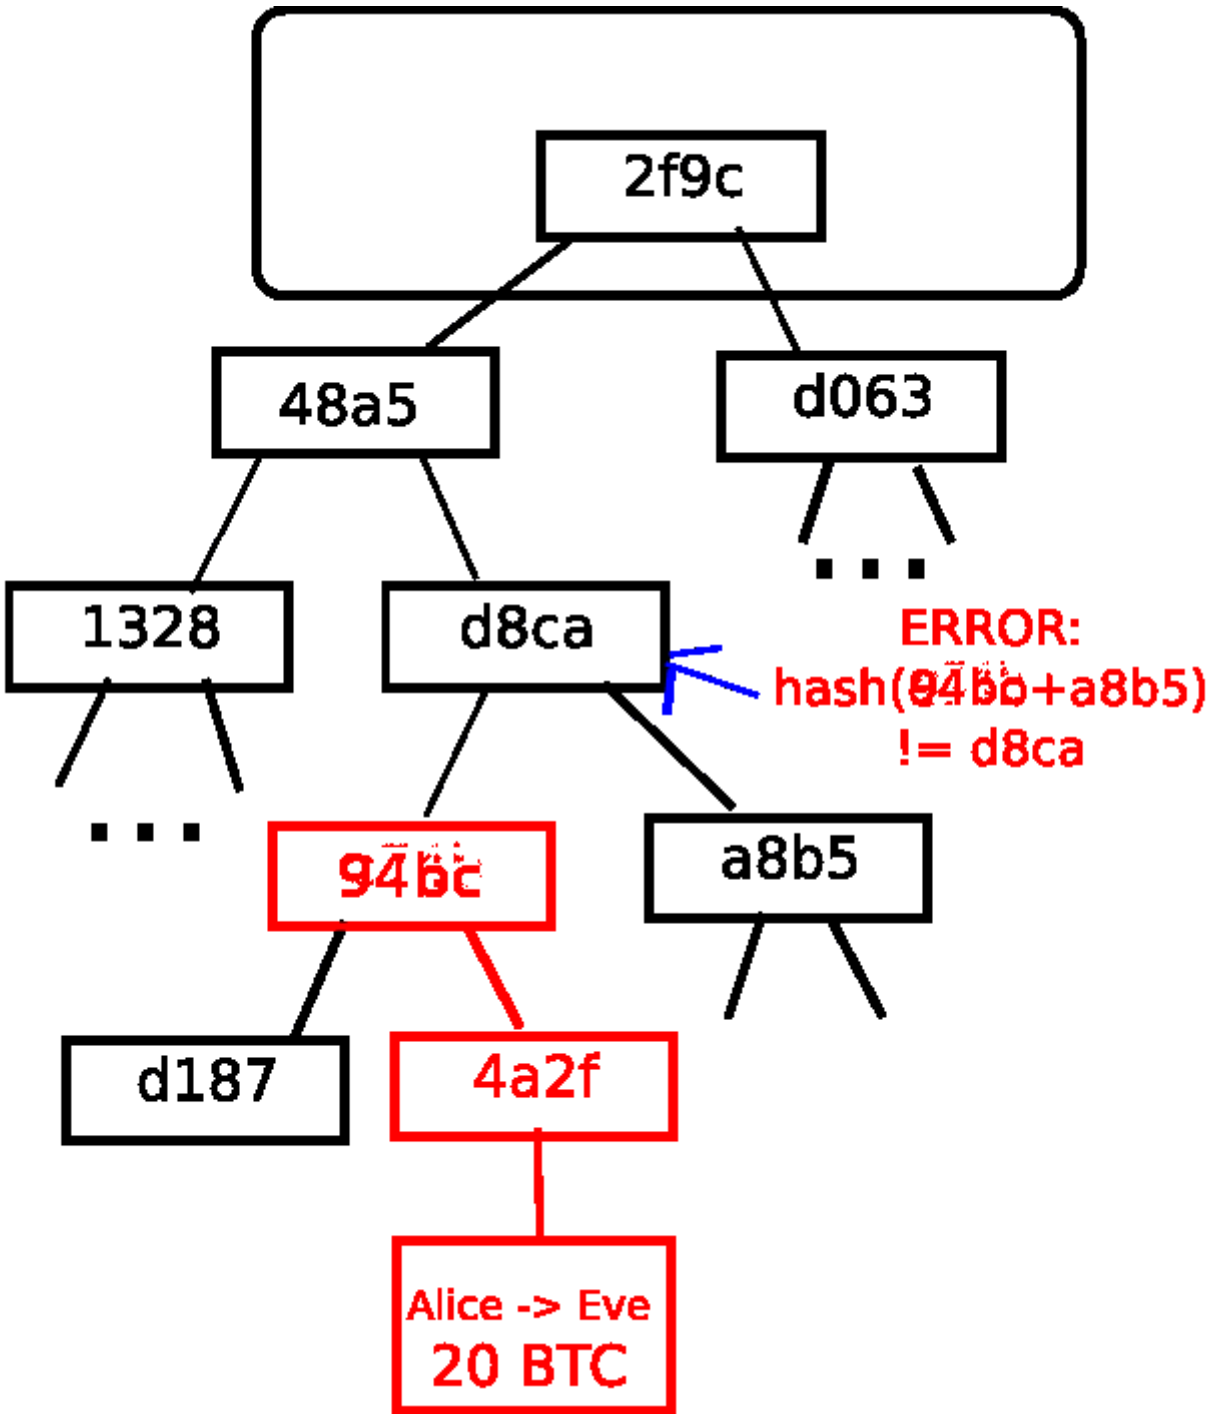


```
02: any attempt to change any part of the Merkle tree will eventually lead
to an inconsistency somewhere up the chain.
```

An important scalability feature of Bitcoin is that the block is stored in a multi-level data structure. The "hash" of a block is actually only the hash of the block header, a roughly 200-byte piece of data that contains the timestamp, nonce, previous block hash and the root hash of a data structure called the Merkle tree storing all transactions in the block. A Merkle tree is a type of binary tree, composed of a set of nodes with a large number of leaf nodes at the bottom of the tree containing the underlying data, a set of intermediate nodes where each node is the hash of its two children, and finally a single root node, also formed from the hash of its two children, representing the "top" of the tree. The purpose of the Merkle tree is to allow the data in a block to be delivered piecemeal: a node can download only the header of a block from one source, the small part of

the tree relevant to them from another source, and still be assured that all of the data is correct. The reason why this works is that hashes propagate upward: if a malicious user attempts to swap in a fake transaction into the bottom of a Merkle tree, this change will cause a change in the node above, and then a change in the node above that, finally changing the root of the tree and therefore the hash of the block, causing the protocol to register it as a completely different block (almost certainly with an invalid proof of work).

The Merkle tree protocol is arguably essential to long-term sustainability. A "full node" in the Bitcoin network, one that stores and processes the entirety of every block, takes up about 15 GB of disk space in the Bitcoin network as of April 2014, and is growing by over a gigabyte per month. Currently, this is viable for some desktop computers and not phones, and later on in the future only businesses and hobbyists will be able to participate. A protocol known as "simplified payment verification" (SPV) allows for another class of nodes to exist, called "light nodes", which download the block headers, verify the proof of work on the block headers, and then download only the "branches" associated with transactions that are relevant to them. This allows light nodes to determine with a strong guarantee of security what the status of any Bitcoin transaction, and their current balance, is while downloading only a very small portion of the entire blockchain.

Alternative Blockchain Applications

The idea of taking the underlying blockchain idea and applying it to other concepts also has a long history. In 2005, Nick Szabo came out with the concept of "secure property titles with owner authority", a document describing how "new advances in replicated database technology" will allow for a blockchain-based system for storing a registry of who owns what land, creating an elaborate framework including concepts such as homesteading, adverse possession and Georgian land tax. However, there was unfortunately no effective replicated database system available at the time, and so the protocol was never implemented in practice. After 2009, however, once Bitcoin's decentralized consensus was developed a number of alternative applications rapidly began to emerge.

Namecoin - created in 2010, Namecoin is best described as a decentralized name registration database. In decentralized protocols like Tor, Bitcoin and BitMessage, there needs to be some way of identifying accounts so that other people can interact with them, but in all existing solutions the only kind of identifier available is a pseudorandom hash like 1LW79wp5ZBqaHW1jL5TCiBCrhQYtHagUWy. Ideally, one would like to be able to have an account with a name like "george". However, the problem is that if one person can create an account named "george" then someone else can use the same process to register "george" for themselves as well and impersonate them. The only solution is a first-to-file paradigm, where the first registerer succeeds and the second fails - a problem perfectly suited for the Bitcoin consensus protocol. Namecoin is the oldest, and most successful, implementation of a name registration system using such an idea.

Colored coins - the purpose of colored coins is to serve as a protocol to allow people to create their own digital currencies - or, in the important trivial case of a currency with one unit, digital tokens, on the Bitcoin blockchain. In the colored coins protocol, one "issues" a new currency by publicly assigning a color to a specific Bitcoin UTXO, and the protocol recursively defines the color of other UTXO to be the same as the color of the inputs that the transaction creating them spent (some special rules apply in the case of mixed-color inputs). This allows users to maintain wallets containing only UTXO of a specific color and send them around much like regular bitcoins, backtracking through the blockchain to determine the color of any UTXO that they receive.

Metacoins - the idea behind a metacoin is to have a protocol that lives on top of Bitcoin, using Bitcoin transactions to store metacoin transactions but having a different state transition function, APPLY'. Because the metacoin protocol cannot prevent invalid metacoin transactions from appearing in the Bitcoin blockchain, a rule is added that if APPLY'(S,TX) returns an error, the protocol defaults to APPLY'(S,TX) = S. This provides an easy mechanism for creating an arbitrary cryptocurrency protocol, potentially with advanced features that cannot be implemented inside of Bitcoin itself, but with a very low development cost since the complexities of mining and networking are already handled by the Bitcoin protocol. Metacoins have been used to implement some classes of financial contracts, name registration and decentralized exchange.

Thus, in general, there are two approaches toward building a consensus protocol: building an independent network, and building a protocol on top of Bitcoin. The former approach, while reasonably successful in the case of applications like Namecoin, is difficult to implement; each individual implementation needs to bootstrap an independent blockchain, as well as building and testing all of the necessary state transition and networking code. Additionally, we predict that the set of applications for decentralized consensus technology will follow a power law distribution where the vast majority of applications would be too small to warrant their own blockchain, and we note that there exist large classes of decentralized applications, particularly decentralized autonomous organizations, that need to interact with each other.
The Bitcoin-based approach, on the other hand, has the flaw that it does not inherit the simplified payment verification features of Bitcoin. SPV works for Bitcoin because it can use blockchain depth as a proxy for validity; at some point, once the ancestors of a transaction go far enough back, it is safe to say that they were legitimately part of the state. Blockchain-based meta-protocols, on the other hand, cannot force the blockchain not to include transactions that are not valid within the context of their own protocols. Hence, a fully secure SPV meta-protocol implementation would need to backward scan all the way to the beginning of the Bitcoin blockchain to determine whether or not certain transactions are valid. Currently, all "light" implementations of Bitcoin-based meta-protocols rely on a trusted server to provide the data, arguably a highly suboptimal result especially when one of the primary purposes of a cryptocurrency is to eliminate the need for trust.

Scripting

Even without any extensions, the Bitcoin protocol actually does facilitate a weak version of a concept of "smart contracts". UTXO in Bitcoin can be owned not just by a public key, but also by a more complicated script expressed in a simple stack-based programming language. In this paradigm, a transaction spending that UTXO must provide data that satisfies the script. Indeed, even the basic public key ownership mechanism is implemented via a script: the script takes an elliptic curve signature as input, verifies it against the transaction and the address that owns the UTXO, and returns 1 if the verification is successful and 0 otherwise. Other, more complicated, scripts exist for various additional use cases. For example, one can construct a script that requires signatures from two out of a given three private keys to validate ("multisig"), a setup useful for corporate accounts, secure savings accounts and some merchant escrow situations. Scripts can also be used to pay bounties for solutions to computational problems, and one can even construct a script that says something like "this Bitcoin UTXO is yours if you can provide an SPV proof that you sent a Dogecoin transaction of this denomination to me", essentially allowing decentralized cross-cryptocurrency exchange.

However, the scripting language as implemented in Bitcoin has several important limitations:

Lack of Turing-completeness - that is to say, while there is a large subset of computation that the Bitcoin

scripting language supports, it does not nearly support everything. The main category that is missing is loops. This is done to avoid infinite loops during transaction verification; theoretically it is a surmountable obstacle for script programmers, since any loop can be simulated by simply repeating the underlying code many times with an if statement, but it does lead to scripts that are very space-inefficient. For example, implementing an alternative elliptic curve signature algorithm would likely require 256 repeated multiplication rounds all individually included in the code.

Value-blindness - there is no way for a UTXO script to provide fine-grained control over the amount that can be withdrawn. For example, one powerful use case of an oracle contract would be a hedging contract, where A and B put in $1000 worth of BTC and after 30 days the script sends $1000 worth of BTC to A and the rest to B. This would require an oracle to determine the value of 1 BTC in USD, but even then it is a massive improvement in terms of trust and infrastructure requirement over the fully centralized solutions that are available now. However, because UTXO are all-or-nothing, the only way to achieve this is through the very inefficient hack of having many UTXO of varying denominations (eg. one UTXO of 2^k for every k up to 30) and having O pick which UTXO to send to A and which to B.

Lack of state - UTXO can either be spent or unspent; there is no opportunity for multi-stage contracts or scripts which keep any other internal state beyond that. This makes it hard to make multi-stage options contracts, decentralized exchange offers or two-stage cryptographic commitment protocols (necessary for secure computational bounties). It also means that UTXO can only be used to build simple, one-off contracts and not more complex "stateful" contracts such as decentralized organizations, and makes meta-protocols difficult to implement. Binary state combined with value-blindness also mean that another important application, withdrawal limits, is impossible.

Blockchain-blindness - UTXO are blind to certain blockchain data such as the nonce and previous block hash. This severely limits applications in gambling, and several other categories, by depriving the scripting language of a potentially valuable source of randomness.

Thus, we see three approaches to building advanced applications on top of cryptocurrency: building a new blockchain, using scripting on top of Bitcoin, and building a meta-protocol on top of Bitcoin. Building a new blockchain allows for unlimited freedom in building a feature set, but at the cost of development time, bootstrapping effort and security. Using scripting is easy to implement and standardize, but is very limited in its capabilities, and meta-protocols, while easy, suffer from faults in scalability. With zkFUND, we intend to build an alternative framework that provides even larger gains in ease of development as well as even stronger light client properties, while at the same time allowing applications to share an economic environment and blockchain security.

zkFUND Accounts

In zkFUND, the state is made up of objects called "accounts", with each account having a 20-byte address and state transitions being direct transfers of value and information between accounts. An zkFUND account contains four fields:

* The nonce, a counter used to make sure each transaction can only be processed once
* The account's current ESPAcoin balance

* The account's contract code, if present
* The account's storage (empty by default)

"ESPAcoin" is the main internal crypto-fuel of zkFUND, and is used to pay transaction fees. In general, there are two types of accounts: externally owned accounts, controlled by private keys, and contract accounts, controlled by their contract code. An externally owned account has no code, and one can send messages from an externally owned account by creating and signing a transaction; in a contract account, every time the contract account receives a message its code activates, allowing it to read and write to internal storage and send other messages or create contracts in turn.
Note that "contracts" in zkFUND should not be seen as something that should be "fulfilled" or "complied with"; rather, they are more like "autonomous agents" that live inside of the zkFUND execution environment, always executing a specific piece of code when "poked" by a message or transaction, and having direct control over their own ESPAcoin balance and their own key/value store to keep track of persistent variables.

Messages and Transactions

The term "transaction" is used in zkFUND to refer to the signed data package that stores a message to be sent from an externally owned account. Transactions contain:

* The recipient of the message
* A signature identifying the sender
* The amount of ESPAcoin to transfer from the sender to the recipient
* An optional data field
* A STARTGAS value, representing the maximum number of computational steps the transaction execution is allowed to take
* A GASPRICE value, representing the fee the sender pays per computational step

The first three are standard fields expected in any cryptocurrency. The data field has no function by default, but the virtual machine has an opcode with which a contract can access the data; as an example use case, if a contract is functioning as an on-blockchain domain registration service, then it may wish to interpret the data being passed to it as containing two "fields", the first field being a domain to register and the second field being the IP address to register it to. The contract would read these values from the message data and appropriately place them in storage.
The STARTGAS and GASPRICE fields are crucial for zkFUND's anti-denial of service model. In order to prevent accidental or hostile infinite loops or other computational wastage in code, each transaction is required to set a limit to how many computational steps of code execution it can use. The fundamental unit of computation is "gas"; usually, a computational step costs 1 gas, but some operations cost higher amounts of gas because they are more computationally expensive, or increase the amount of data that must be stored as part of the state. There is also a fee of 5 gas for every byte in the transaction data. The intent of the fee system is to require an attacker to pay proportionately for every resource that they consume, including computation, bandwidth and storage; hence, any transaction that leads to the network consuming a greater amount of any of these resources must have a gas fee roughly proportional to the increment.

Contracts have the ability to send "messages" to other contracts. Messages are virtual objects that are never serialized and exist only in the zkFUND execution environment. A message contains:

* The sender of the message (implicit)
* The recipient of the message
* The amount of ESPAcoin to transfer alongside the message
* An optional data field
* A STARTGAS value

Essentially, a message is like a transaction, except it is produced by a contract and not an external actor. A message is produced when a contract currently executing code executes the CALL opcode, which produces and executes a message. Like a transaction, a message leads to the recipient account running its code. Thus, contracts can have relationships with other contracts in exactly the same way that external actors can.
Note that the gas allowance assigned by a transaction or contract applies to the total gas consumed by that transaction and all sub-executions. For example, if an external actor A sends a transaction to B with 1000 gas, and B consumes 600 gas before sending a message to C, and the internal execution of C consumes 300 gas before returning, then B can spend another 100 gas before running out of gas.

zkFUND State Transition Function

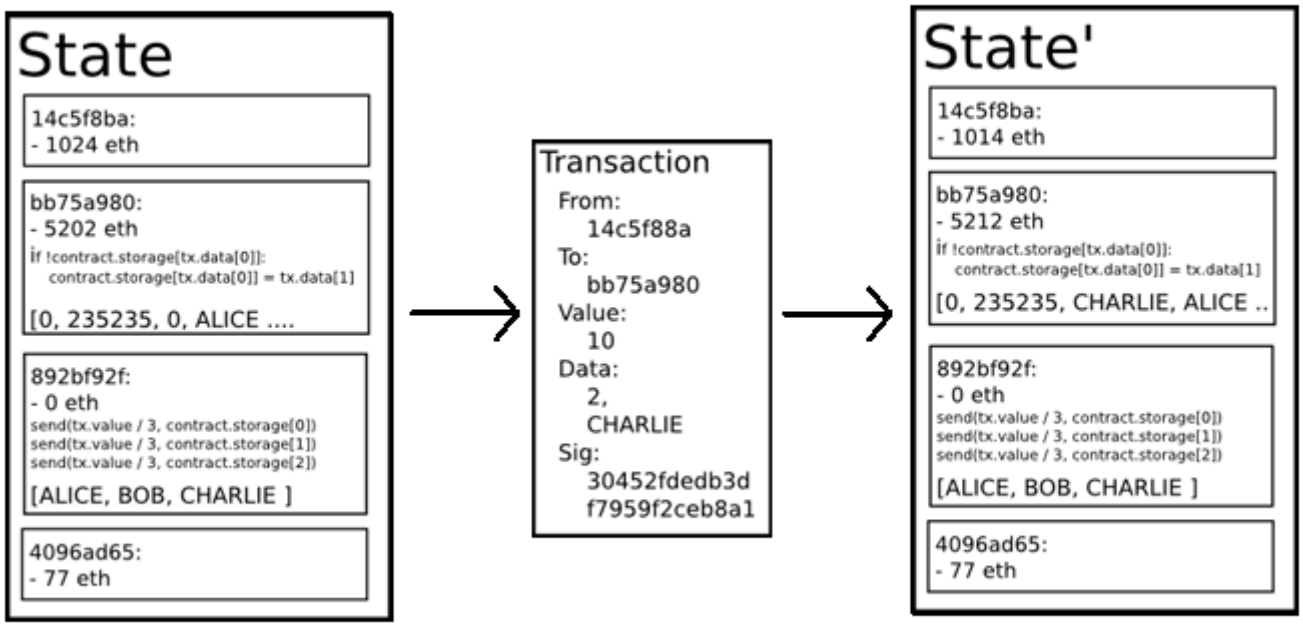

The zkFUND state transition function, APPLY(S,TX) -> S' can be defined as follows:

1. Check if the transaction is well-formed (ie. has the right number of values), the signature is valid, and the nonce matches the nonce in the sender's account. If not, return an error.
2. Calculate the transaction fee as STARTGAS * GASPRICE, and determine the sending address from the signature. Subtract the fee from the sender's account balance and increment the sender's nonce. If there is not enough balance to spend, return an error.
3. Initialize GAS = STARTGAS, and take off a certain quantity of gas per byte to pay for the bytes in the

transaction.
4. Transfer the transaction value from the sender's account to the receiving account. If the receiving account does not yet exist, create it. If the receiving account is a contract, run the contract's code either to completion or until the execution runs out of gas.
5. If the value transfer failed because the sender did not have enough money, or the code execution ran out of gas, revert all state changes except the payment of the fees, and add the fees to the miner's account.
6. Otherwise, refund the fees for all remaining gas to the sender, and send the fees paid for gas consumed to the miner.

For example, suppose that the contract's code is:

```
if !self.storage[calldataload(0)]:
                        self.storage[calldataload(0)] = calldataload(32)
```

Note that in reality the contract code is written in the low-level EVM code; this example is written in Serpent, one of our high-level languages, for clarity, and can be compiled down to EVM code. Suppose that the contract's storage starts off empty, and a transaction is sent with 10 ESPAcoin value, 2000 gas, 0.001 ESPAcoin gasprice, and 64 bytes of data, with bytes 0-31 representing the number 2 and bytes 32-63 representing the string CHARLIE. The process for the state transition function in this case is as follows:

1. Check that the transaction is valid and well formed.
2. Check that the transaction sender has at least 2000 * 0.001 = 2 ESPAcoin. If it is, then subtract 2 ESPAcoin from the sender's account.
3. Initialize gas = 2000; assuming the transaction is 170 bytes long and the byte-fee is 5, subtract 850 so that there is 1150 gas left.
4. Subtract 10 more ESPAcoin from the sender's account, and add it to the contract's account.
5. Run the code. In this case, this is simple: it checks if the contract's storage at index 2 is used, notices that it is not, and so it sets the storage at index 2 to the value CHARLIE. Suppose this takes 187 gas, so the remaining amount of gas is 1150 - 187 = 963
6. Add 963 * 0.001 = 0.963 ESPAcoin back to the sender's account, and return the resulting state.

If there was no contract at the receiving end of the transaction, then the total transaction fee would simply be equal to the provided GASPRICE multiplied by the length of the transaction in bytes, and the data sent alongside the transaction would be irrelevant.
Note that messages work equivalently to transactions in terms of reverts: if a message execution runs out of gas, then that message's execution, and all other executions triggered by that execution, revert, but parent executions do not need to revert. This means that it is "safe" for a contract to call another contract, as if A calls B with G gas then A's execution is guaranteed to lose at most G gas. Finally, note that there is an opcode, CREATE, that creates a contract; its execution mechanics are generally similar to CALL, with the exception that the output of the execution determines the code of a newly created contract.

Code Execution

The code in zkFUND contracts is written in a low-level, stack-based bytecode language, referred to as "zkFUND virtual machine code" or "EVM code". The code consists of a series of bytes, where each byte

represents an operation. In general, code execution is an infinite loop that consists of repeatedly carrying out the operation at the current program counter (which begins at zero) and then incrementing the program counter by one, until the end of the code is reached or an error or STOP or RETURN instruction is detected. The operations have access to three types of space in which to store data:

* The stack, a last-in-first-out container to which values can be pushed and popped
* Memory, an infinitely expandable byte array
* The contract's long-term storage, a key/value store. Unlike stack and memory, which reset after computation ends, storage persists for the long term.

The code can also access the value, sender and data of the incoming message, as well as block header data, and the code can also return a byte array of data as an output.
The formal execution model of EVM code is surprisingly simple. While the zkFUND virtual machine is running, its full computational state can be defined by the tuple (block_state, transaction, message, code, memory, stack, pc, gas), where block_state is the global state containing all accounts and includes balances and storage. At the start of every round of execution, the current instruction is found by taking the pcth (Program Counter) byte of code (or 0 if pc >= len(code)), and each instruction has its own definition in terms of how it affects the tuple. For example, ADD pops two items off the stack and pushes their sum, reduces gas by 1 and increments pc by 1, and SSTORE pops the top two items off the stack and inserts the second item into the contract's storage at the index specified by the first item. Although there are many ways to optimize zkFUND virtual machine execution via just-in-time compilation, a basic implementation of zkFUND can be done in a few hundred lines of code.

Blockchain and Mining

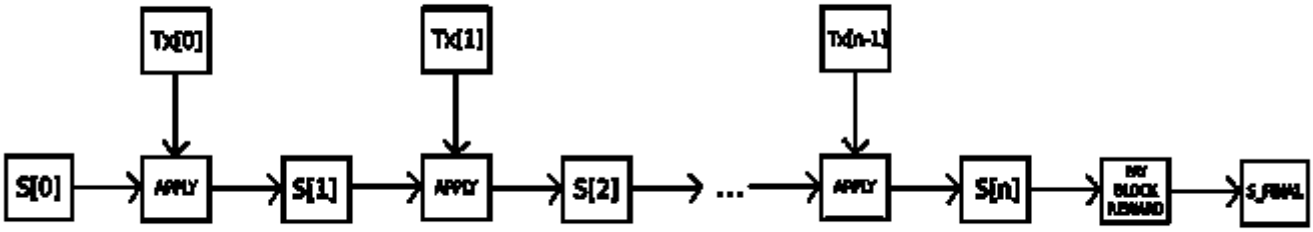

The zkFUND blockchain is in many ways similar to the Bitcoin blockchain, although it does have some differences. The main difference between zkFUND and Bitcoin with regard to the blockchain architecture is that, unlike Bitcoin, zkFUND blocks contain a copy of both the transaction list and the most recent state. Aside from that, two other values, the block number and the difficulty, are also stored in the block. The basic block validation algorithm in zkFUND is as follows:

1. Check if the previous block referenced exists and is valid.
2. Check that the timestamp of the block is greater than that of the referenced previous block and less than 15 minutes into the future
3. Check that the block number, difficulty, transaction root, uncle root and gas limit (various low-level

zkFUND-specific concepts) are valid.
4. Check that the proof of work on the block is valid.
5. Let S[0] be the state at the end of the previous block.
6. Let TX be the block's transaction list, with n transactions. For all i in 0...n-1, set S[i+1] = APPLY(S[i],TX[i]). If any application returns an error, or if the total gas consumed in the block up until this point exceeds the GASLIMIT, return an error.
7. Let S_FINAL be S[n], but adding the block reward paid to the miner.
8. Check if the Merkle tree root of the state S_FINAL is equal to the final state root provided in the block header. If it is, the block is valid; otherwise, it is not valid.

The approach may seem highly inefficient at first glance, because it needs to store the entire state with each block, but in reality efficiency should be comparable to that of Bitcoin. The reason is that the state is stored in the tree structure, and after every block only a small part of the tree needs to be changed. Thus, in general, between two adjacent blocks the vast majority of the tree should be the same, and therefore the data can be stored once and referenced twice using pointers (ie. hashes of subtrees). A special kind of tree known as a "Patricia tree" is used to accomplish this, including a modification to the Merkle tree concept that allows for nodes to be inserted and deleted, and not just changed, efficiently. Additionally, because all of the state information is part of the last block, there is no need to store the entire blockchain history - a strategy which, if it could be applied to Bitcoin, can be calculated to provide 5-20x savings in space.
A commonly asked question is "where" contract code is executed, in terms of physical hardware. This has a simple answer: the process of executing contract code is part of the definition of the state transition function, which is part of the block validation algorithm, so if a transaction is added into block B the code execution spawned by that transaction will be executed by all nodes, now and in the future, that download and validate block B.

Token Systems

In general, there are three types of applications on top of zkFUND. The first category is financial applications, providing users with more powerful ways of managing and entering into contracts using their money. This includes sub-currencies, financial derivatives, hedging contracts, savings wallets, wills, and ultimately even some classes of full-scale employment contracts. The second category is semi-financial applications, where money is involved but there is also a heavy non-monetary side to what is being done; a perfect example is self-enforcing bounties for solutions to computational problems. Finally, there are applications such as online voting and decentralized governance that are not financial at all.
On-blockchain token systems have many applications ranging from sub-currencies representing assets such as USD or gold to company stocks, individual tokens representing smart property, secure unforgeable coupons, and even token systems with no ties to conventional value at all, used as point systems for incentivization. Token systems are surprisingly easy to implement in zkFUND. The key point to understand is that all a currency, or token system, fundamentally is a database with one operation: subtract X units from A and give X units to B, with the proviso that (i) A had at least X units before the transaction and (2) the transaction is approved by A. All that it takes to implement a token system is to implement this logic into a contract.
The basic code for implementing a token system in Serpent looks as follows:

```
def send(to, value):
        if self.storage[msg.sender] >= value:
```

```
                self.storage[msg.sender] = self.storage[msg.sender] - value
                self.storage[to] = self.storage[to] + value
```

This is essentially a literal implementation of the "banking system" state transition function described further above in this document. A few extra lines of code need to be added to provide for the initial step of distributing the currency units in the first place and a few other edge cases, and ideally a function would be added to let other contracts query for the balance of an address. But that's all there is to it. Theoretically, zkFUND-based token systems acting as sub-currencies can potentially include another important feature that on-chain Bitcoin-based meta-currencies lack: the ability to pay transaction fees directly in that currency. The way this would be implemented is that the contract would maintain an ESPAcoin balance with which it would refund ESPAcoin used to pay fees to the sender, and it would refill this balance by collecting the internal currency units that it takes in fees and reselling them in a constant running auction. Users would thus need to "activate" their accounts with ESPAcoin, but once the ESPAcoin is there it would be reusable because the contract would refund it each time.

Financial derivatives

Financial derivatives are the most common application of a "smart contract", and one of the simplest to implement in code. The main challenge in implementing financial contracts is that the majority of them require reference to an external price ticker; for example, a very desirable application is a smart contract that hedges against the volatility of ESPAcoin (or another cryptocurrency) with respect to the US dollar, but doing this requires the contract to know what the value of ESPAcoin/USD is. The simplest way to do this is through a "data feed" contract maintained by a specific party (eg. NASDAQ) designed so that that party has the ability to update the contract as needed, and providing an interface that allows other contracts to send a message to that contract and get back a response that provides the price.

Given that critical ingredient, the hedging contract would look as follows:

1. Wait for party A to input 1000 ESPAcoin.
2. Wait for party B to input 1000 ESPAcoin.
3. Record the USD value of 1000 ESPAcoin, calculated by querying the data feed contract, in storage, say this is $x.
4. After 30 days, allow A or B to "reactivate" the contract in order to send $x worth of ESPAcoin (calculated by querying the data feed contract again to get the new price) to A and the rest to B.

Such a contract would have significant potential in crypto-commerce. One of the main problems cited about cryptocurrency is the fact that it's volatile; although many users and merchants may want the security and convenience of dealing with cryptographic assets, they may not wish to face that prospect of losing 23% of the value of their funds in a single day. Up until now, the most commonly proposed solution has been issuer-backed assets; the idea is that an issuer creates a sub-currency in which they have the right to issue and revoke units, and provide one unit of the currency to anyone who provides them (offline) with one unit of a specified underlying asset (eg. gold, USD). The issuer then promises to provide one unit of the underlying asset to anyone who sends back one unit of the crypto-asset. This mechanism allows any non-cryptographic asset to be "uplifted" into a cryptographic asset, provided that the issuer can be trusted.
In practice, however, issuers are not always trustworthy, and in some cases the banking infrastructure is too

weak, or too hostile, for such services to exist. Financial derivatives provide an alternative. Here, instead of a single issuer providing the funds to back up an asset, a decentralized market of speculators, betting that the price of a cryptographic reference asset (eg. ESPAcoin) will go up, plays that role. Unlike issuers, speculators have no option to default on their side of the bargain because the hedging contract holds their funds in escrow. Note that this approach is not fully decentralized, because a trusted source is still needed to provide the price ticker, although arguably even still this is a massive improvement in terms of reducing infrastructure requirements (unlike being an issuer, issuing a price feed requires no licenses and can likely be categorized as free speech) and reducing the potential for fraud.

Identity and Reputation Systems

The earliest alternative cryptocurrency of all, Namecoin, attempted to use a Bitcoin-like blockchain to provide a name registration system, where users can register their names in a public database alongside other data. The major cited use case is for a DNS system, mapping domain names like "bitcoin.org" (or, in Namecoin's case, "bitcoin.bit") to an IP address. Other use cases include email authentication and potentially more advanced reputation systems. Here is the basic contract to provide a Namecoin-like name registration system on zkFUND:

```
def register(name, value):
    if !self.storage[name]:
        self.storage[name] = value
```

The contract is very simple; all it is is a database inside the zkFUND network that can be added to, but not modified or removed from. Anyone can register a name with some value, and that registration then sticks forever. A more sophisticated name registration contract will also have a "function clause" allowing other contracts to query it, as well as a mechanism for the "owner" (ie. the first registerer) of a name to change the data or transfer ownership. One can even add reputation and web-of-trust functionality on top.

Decentralized File Storage

Over the past few years, there have emerged a number of popular online file storage startups, the most prominent being Dropbox, seeking to allow users to upload a backup of their hard drive and have the service store the backup and allow the user to access it in exchange for a monthly fee. However, at this point the file storage market is at times relatively inefficient; a cursory look at various existing solutions shows that, particularly at the "uncanny valley" 20-200 GB level at which neither free quotas nor enterprise-level discounts kick in, monthly prices for mainstream file storage costs are such that you are paying for more than the cost of the entire hard drive in a single month. zkFUND contracts can allow for the development of a decentralized file storage ecosystem, where individual users can earn small quantities of money by renting out their own hard drives and unused space can be used to further drive down the costs of file storage.

The key underpinning piece of such a device would be what we have termed the "decentralized Dropbox contract". This contract works as follows. First, one splits the desired data up into blocks, encrypting each block for privacy, and builds a Merkle tree out of it. One then makes a contract with the rule that, every N blocks, the contract would pick a random index in the Merkle tree (using the previous block hash, accessible from contract code, as a source of randomness), and give X ESPAcoin to the first entity to supply a transaction with a simplified payment verification-like proof of ownership of the block at that particular index in the tree. When a user wants to re-download their file, they can use a micropayment channel protocol (eg. pay 1 szabo

per 32 kilobytes) to recover the file; the most fee-efficient approach is for the payer not to publish the transaction until the end, instead replacing the transaction with a slightly more lucrative one with the same nonce after every 32 kilobytes.

An important feature of the protocol is that, although it may seem like one is trusting many random nodes not to decide to forget the file, one can reduce that risk down to near-zero by splitting the file into many pieces via secret sharing, and watching the contracts to see each piece is still in some node's possession. If a contract is still paying out money, that provides a cryptographic proof that someone out there is still storing the file.

Decentralized Autonomous Organizations

The general concept of a "decentralized autonomous organization" is that of a virtual entity that has a certain set of members or shareholders which, perhaps with a 67% majority, have the right to spend the entity's funds and modify its code. The members would collectively decide on how the organization should allocate its funds. Methods for allocating a DAO's funds could range from bounties, salaries to even more exotic mechanisms such as an internal currency to reward work. This essentially replicates the legal trappings of a traditional company or nonprofit but using only cryptographic blockchain technology for enforcement. So far much of the talk around DAOs has been around the "capitalist" model of a "decentralized autonomous corporation" (DAC) with dividend-receiving shareholders and tradable shares; an alternative, perhaps described as a "decentralized autonomous community", would have all members have an equal share in the decision making and require 67% of existing members to agree to add or remove a member. The requirement that one person can only have one membership would then need to be enforced collectively by the group.

A general outline for how to code a DAO is as follows. The simplest design is simply a piece of self-modifying code that changes if two thirds of members agree on a change. Although code is theoretically immutable, one can easily get around this and have de-facto mutability by having chunks of the code in separate contracts, and having the address of which contracts to call stored in the modifiable storage. In a simple implementation of such a DAO contract, there would be three transaction types, distinquished by the data provided in the transaction:

* [0,i,K,V] to register a proposal with index i to change the address at storage index K to value V
* [0,i] to register a vote in favor of proposal i
* [2,i] to finalize proposal i if enough votes have been made

The contract would then have clauses for each of these. It would maintain a record of all open storage changes, along with a list of who voted for them. It would also have a list of all members. When any storage change gets to two thirds of members voting for it, a finalizing transaction could execute the change. A more sophisticated skeleton would also have built-in voting ability for features like sending a transaction, adding members and removing members, and may even provide for Liquid Democracy-style vote delegation (ie. anyone can assign someone to vote for them, and assignment is transitive so if A assigns B and B assigns C then C determines A's vote). This design would allow the DAO to grow organically as a decentralized community, allowing people to eventually delegate the task of filtering out who is a member to specialists, although unlike in the "current system" specialists can easily pop in and out of existence over time as individual community members change their alignments.

An alternative model is for a decentralized corporation, where any account can have zero or more shares, and two thirds of the shares are required to make a decision. A complete skeleton would involve asset management functionality, the ability to make an offer to buy or sell shares, and the ability to accept offers (preferably with

an order-matching mechanism inside the contract). Delegation would also exist Liquid Democracy-style, generalizing the concept of a "board of directors".

Further Applications

1. Savings wallets. Suppose that Alice wants to keep her funds safe, but is worried that she will lose or someone will hack her private key. She puts ESPAcoin into a contract with Bob, a bank, as follows:

* Alice alone can withdraw a maximum of 1% of the funds per day.
* Bob alone can withdraw a maximum of 1% of the funds per day, but Alice has the ability to make a transaction with her key shutting off this ability.
* Alice and Bob together can withdraw anything.

2. Crop insurance. One can easily make a financial derivatives contract but using a data feed of the weather instead of any price index. If a farmer in Iowa purchases a derivative that pays out inversely based on the precipitation in Iowa, then if there is a drought, the farmer will automatically receive money and if there is enough rain the farmer will be happy because their crops would do well. This can be expanded to natural disaster insurance generally.

3. A decentralized data feed. For financial contracts for difference, it may actually be possible to decentralize the data feed via a protocol called "SchellingCoin". SchellingCoin basically works as follows: N parties all put into the system the value of a given datum (eg. the ESPAcoin/USD price), the values are sorted, and everyone between the 25th and 75th percentile gets one token as a reward. Everyone has the incentive to provide the answer that everyone else will provide, and the only value that a large number of players can realistically agree on is the obvious default: the truth. This creates a decentralized protocol that can theoretically provide any number of values, including the ESPAcoin/USD price, the temperature in Berlin or even the result of a particular hard computation.

4. Smart multisignature escrow. Bitcoin allows multisignature transaction contracts where, for example, three out of a given five keys can spend the funds. zkFUND allows for more granularity; for example, four out of five can spend everything, three out of five can spend up to 10% per day, and two out of five can spend up to 0.5% per day. Additionally, zkFUND multisig is asynchronous - two parties can register their signatures on the blockchain at different times and the last signature will automatically send the transaction.

5. Cloud computing. The EVM technology can also be used to create a verifiable computing environment, allowing users to ask others to carry out computations and then optionally ask for proofs that computations at certain randomly selected checkpoints were done correctly. This allows for the creation of a cloud computing market where any user can participate with their desktop, laptop or specialized server, and spot-checking together with security deposits can be used to ensure that the system is trustworthy (ie. nodes cannot profitably cheat). Although such a system may not be suitable for all tasks; tasks that require a high level of inter-process communication, for example, cannot easily be done on a large cloud of nodes. Other tasks, however, are much easier to parallelize; projects like SETI@home, folding@home and genetic algorithms can easily be implemented on top of such a platform.

6. Peer-to-peer gambling. Any number of peer-to-peer gambling protocols, such as Frank Stajano and Richard

Clayton's Cyberdice, can be implemented on the zkFUND blockchain. The simplest gambling protocol is actually simply a contract for difference on the next block hash, and more advanced protocols can be built up from there, creating gambling services with near-zero fees that have no ability to cheat.

7. Prediction markets. Provided an oracle or SchellingCoin, prediction markets are also easy to implement, and prediction markets together with SchellingCoin may prove to be the first mainstream application of futarchy as a governance protocol for decentralized organizations.

8. On-chain decentralized marketplaces, using the identity and reputation system as a base.

Modified GHOST Implementation

The "Greedy Heaviest Observed Subtree" (GHOST) protocol is an innovation first introduced by Yonatan Sompolinsky and Aviv Zohar in December 2013. The motivation behind GHOST is that blockchains with fast confirmation times currently suffer from reduced security due to a high stale rate - because blocks take a certain time to propagate through the network, if miner A mines a block and then miner B happens to mine another block before miner A's block propagates to B, miner B's block will end up wasted and will not contribute to network security. Furthermore, there is a centralization issue: if miner A is a mining pool with 30% hashpower and B has 10% hashpower, A will have a risk of producing a stale block 70% of the time (since the other 30% of the time A produced the last block and so will get mining data immediately) whereas B will have a risk of producing a stale block 90% of the time. Thus, if the block interval is short enough for the stale rate to be high, A will be substantially more efficient simply by virtue of its size. With these two effects combined, blockchains which produce blocks quickly are very likely to lead to one mining pool having a large enough percentage of the network hashpower to have de facto control over the mining process.

As described by Sompolinsky and Zohar, GHOST solves the first issue of network security loss by including stale blocks in the calculation of which chain is the "longest"; that is to say, not just the parent and further ancestors of a block, but also the stale descendants of the block's ancestor (in zkFUND jargon, "uncles") are added to the calculation of which block has the largest total proof of work backing it. To solve the second issue of centralization bias, we go beyond the protocol described by Sompolinsky and Zohar, and also provide block rewards to stales: a stale block receives 87.5% of its base reward, and the nephew that includes the stale block receives the remaining 12.5%. Transaction fees, however, are not awarded to uncles.

zkFUND implements a simplified version of GHOST which only goes down seven levels. Specifically, it is defined as follows:

* A block must specify a parent, and it must specify 0 or more uncles
* An uncle included in block B must have the following properties:
* -> It must be a direct child of the kth generation ancestor of B, where 2 <= k <= 7.
* -> It cannot be an ancestor of B
* -> An uncle must be a valid block header, but does not need to be a previously verified or even valid block
* -> An uncle must be different from all uncles included in previous blocks and all other uncles included in the same block (non-double-inclusion)
* For every uncle U in block B, the miner of B gets an additional 3.125% added to its coinbase reward and the miner of U gets 93.75% of a standard coinbase reward.

This limited version of GHOST, with uncles includable only up to 7 generations, was used for two reasons.

First, unlimited GHOST would include too many complications into the calculation of which uncles for a given block are valid. Second, unlimited GHOST with compensation as used in zkFUND removes the incentive for a miner to mine on the main chain and not the chain of a public attacker.

Fees

Because every transaction published into the blockchain imposes on the network the cost of needing to download and verify it, there is a need for some regulatory mechanism, typically involving transaction fees, to prevent abuse. The default approach, used in Bitcoin, is to have purely voluntary fees, relying on miners to act as the gatekeepers and set dynamic minimums. This approach has been received very favorably in the Bitcoin community particularly because it is "market-based", allowing supply and demand between miners and transaction senders determine the price. The problem with this line of reasoning is, however, that transaction processing is not a market; although it is intuitively attractive to construe transaction processing as a service that the miner is offering to the sender, in reality every transaction that a miner includes will need to be processed by every node in the network, so the vast majority of the cost of transaction processing is borne by third parties and not the miner that is making the decision of whether or not to include it. Hence, tragedy-of-the-commons problems are very likely to occur.
However, as it turns out this flaw in the market-based mechanism, when given a particular inaccurate simplifying assumption, magically cancels itself out. The argument is as follows. Suppose that:

1. A transaction leads to k operations, offering the reward kR to any miner that includes it where R is set by the sender and k and R are (roughly) visible to the miner beforehand.
2. An operation has a processing cost of C to any node (ie. all nodes have equal efficiency)
3. There are N mining nodes, each with exactly equal processing power (ie. 1/N of total)
4. No non-mining full nodes exist.

A miner would be willing to process a transaction if the expected reward is greater than the cost. Thus, the expected reward is kR/N since the miner has a 1/N chance of processing the next block, and the processing cost for the miner is simply kC. Hence, miners will include transactions where kR/N > kC, or R > NC. Note that R is the per-operation fee provided by the sender, and is thus a lower bound on the benefit that the sender derives from the transaction, and NC is the cost to the entire network together of processing an operation. Hence, miners have the incentive to include only those transactions for which the total utilitarian benefit exceeds the cost.
However, there are several important deviations from those assumptions in reality:

1. The miner does pay a higher cost to process the transaction than the other verifying nodes, since the extra verification time delays block propagation and thus increases the chance the block will become a stale.
2. There do exist nonmining full nodes.
3. The mining power distribution may end up radically inegalitarian in practice.
4. Speculators, political enemies and crazies whose utility function includes causing harm to the network do exist, and they can cleverly set up contracts where their cost is much lower than the cost paid by other verifying nodes.

(1) provides a tendency for the miner to include fewer transactions, and (2) increases NC; hence, these two effects at least partially cancel each other out. (3) and (4) are the major issue; to solve them we simply institute

a floating cap: no block can have more operations than BLK_LIMIT_FACTOR times the long-term exponential moving average. Specifically:

```
blk.oplimit = floor((blk.parent.oplimit * (EMAFACTOR - 1) +
floor(parent.opcount * BLK_LIMIT_FACTOR)) / EMA_FACTOR)
```

BLK_LIMIT_FACTOR and EMA_FACTOR are constants that will be set to 65536 and 1.5 for the time being, but will likely be changed after further analysis.
There is another factor disincentivizing large block sizes in Bitcoin: blocks that are large will take longer to propagate, and thus have a higher probability of becoming stales. In zkFUND, highly gas-consuming blocks can also take longer to propagate both because they are physically larger and because they take longer to process the transaction state transitions to validate. This delay disincentive is a significant consideration in Bitcoin, but less so in zkFUND because of the GHOST protocol; hence, relying on regulated block limits provides a more stable baseline.
zkFUND Network fees, which differ from blockchain fees, are paid directly between participants within the channel. The fees pay for the time-value of money for consuming the channel for a determined maximum period of time, and for counterparty risk of non-communication.
Counterparty risk for fees only exist with one's direct channel counterparty. If a node two hops away decides to disconnect and their transaction gets broadcast on the blockchain, one's direct counterparties should not broadcast on the blockchain, but continue to update via novation with a new Commitment Transaction. See the Decrementing Timelocks entry in the HTLC section for more information about counterparty risk.
The time-value of fees pays for consuming time (e.g. 3 days) and is conceptually equivalent to a gold lease rate without custodial risk; it is the time-value for using up the access to money for a very short duration. Since certain paths may become very profitable in one direction, it is possible for fees to be negative to encourage the channel to be available for those profitable paths.

Computation And Turing-Completeness

An important note is that the zkFUND virtual machine is Turing-complete; this means that EVM code can encode any computation that can be conceivably carried out, including infinite loops. EVM code allows looping in two ways. First, there is a JUMP instruction that allows the program to jump back to a previous spot in the code, and a JUMPI instruction to do conditional jumping, allowing for statements like while x < 27: x = x * 2. Second, contracts can call other contracts, potentially allowing for looping through recursion. This naturally leads to a problem: can malicious users essentially shut miners and full nodes down by forcing them to enter into an infinite loop? The issue arises because of a problem in computer science known as the halting problem: there is no way to tell, in the general case, whether or not a given program will ever halt.
As described in the state transition section, our solution works by requiring a transaction to set a maximum number of computational steps that it is allowed to take, and if execution takes longer computation is reverted but fees are still paid. Messages work in the same way. To show the motivation behind our solution, consider the following examples:

* An attacker creates a contract which runs an infinite loop, and then sends a transaction activating that loop to the miner. The miner will process the transaction, running the infinite loop, and wait for it to run out of gas. Even though the execution runs out of gas and stops halfway through, the transaction is still valid and the miner still claims the fee from the attacker for each computational step.

* An attacker creates a very long infinite loop with the intent of forcing the miner to keep computing for such a long time that by the time computation finishes a few more blocks will have come out and it will not be possible for the miner to include the transaction to claim the fee. However, the attacker will be required to submit a value for STARTGAS limiting the number of computational steps that execution can take, so the miner will know ahead of time that the computation will take an excessively large number of steps.

* An attacker sees a contract with code of some form like send(A,contract.storage[A]); contract.storage[A] = 0, and sends a transaction with just enough gas to run the first step but not the second (ie. making a withdrawal but not letting the balance go down). The contract author does not need to worry about protecting against such attacks, because if execution stops halfway through the changes get reverted.

* A financial contract works by taking the median of nine proprietary data feeds in order to minimize risk. An attacker takes over one of the data feeds, which is designed to be modifiable via the variable-address-call mechanism described in the section on DAOs, and converts it to run an infinite loop, thereby attempting to force any attempts to claim funds from the financial contract to run out of gas. However, the financial contract can set a gas limit on the message to prevent this problem.

The alternative to Turing-completeness is Turing-incompleteness, where JUMP and JUMPI do not exist and only one copy of each contract is allowed to exist in the call stack at any given time. With this system, the fee system described and the uncertainties around the effectiveness of our solution might not be necessary, as the cost of executing a contract would be bounded above by its size. Additionally, Turing-incompleteness is not even that big a limitation; out of all the contract examples we have conceived internally, so far only one required a loop, and even that loop could be removed by making 26 repetitions of a one-line piece of code. Given the serious implications of Turing-completeness, and the limited benefit, why not simply have a Turing-incomplete language? In reality, however, Turing-incompleteness is far from a neat solution to the problem. To see why, consider the following contracts:

```
C0: call(C1); call(C1);
C1: call(C2); call(C2);
C2: call(C3); call(C3);
...
C49: call(C50); call(C50);
C50: (run one step of a program and record the change in storage)
```

Now, send a transaction to A. Thus, in 51 transactions, we have a contract that takes up 250 computational steps. Miners could try to detect such logic bombs ahead of time by maintaining a value alongside each contract specifying the maximum number of computational steps that it can take, and calculating this for contracts calling other contracts recursively, but that would require miners to forbid contracts that create other contracts (since the creation and execution of all 26 contracts above could easily be rolled into a single contract). Another problematic point is that the address field of a message is a variable, so in general it may not even be possible to tell which other contracts a given contract will call ahead of time. Hence, all in all, we have a surprising conclusion: Turing-completeness is surprisingly easy to manage, and the lack of Turing-completeness is equally surprisingly difficult to manage unless the exact same controls are in place -

but in that case why not just let the protocol be Turing-complete?

Currency And Issuance

The zkFUND network includes its own built-in currency, ESPAcoin, which serves the dual purpose of providing a primary liquidity layer to allow for efficient exchange between various types of digital assets and, more importantly, of providing a mechanism for paying transaction fees. For convenience and to avoid future argument (see the current mBTC/uBTC/satoshi debate in Bitcoin), the denominations will be pre-labeled:

* 1: wei
* 10^12: szabo
* 10^15: finney
* 10^18: ESPAcoin

This should be taken as an expanded version of the concept of "dollars" and "cents" or "BTC" and "satoshi". In the near future, we expect "ESPAcoin" to be used for ordinary transactions, "finney" for microtransactions and "szabo" and "wei" for technical discussions around fees and protocol implementation; the remaining denominations may become useful later and should not be included in clients at this point.
The issuance model will be as follows:

* ESPAcoin will be released in a currency sale at the price of 1000-2000 ESPAcoin per BTC, a mechanism intended to fund the zkFUND organization and pay for development that has been used with success by other platforms such as Mastercoin and NXT. Earlier buyers will benefit from larger discounts. The BTC received from the sale will be used entirely to pay salaries and bounties to developers and invested into various for-profit and non-profit projects in the zkFUND and cryptocurrency ecosystem.

* 0.099x the total amount sold (60102216 ESPAcoin) will be allocated to the organization to compensate early contributors and pay ESPAcoin-denominated expenses before the genesis block.

* 0.099x the total amount sold will be maintained as a long-term reserve.

* 0.26x the total amount sold will be allocated to miners per year forever after that point.

Group	At launch	After 1 year	After 5 years
Currency units	1.198X	1.458X	2.498X
Purchasers	83.5%	68.6%	40.0%
Reserve spent pre-sale	8.26%	6.79%	3.96%
Reserve used post-sale	8.26%	6.79%	3.96%
Miners	0%	17.8%	52.0%

Long-Term Supply Growth Rate (percent)

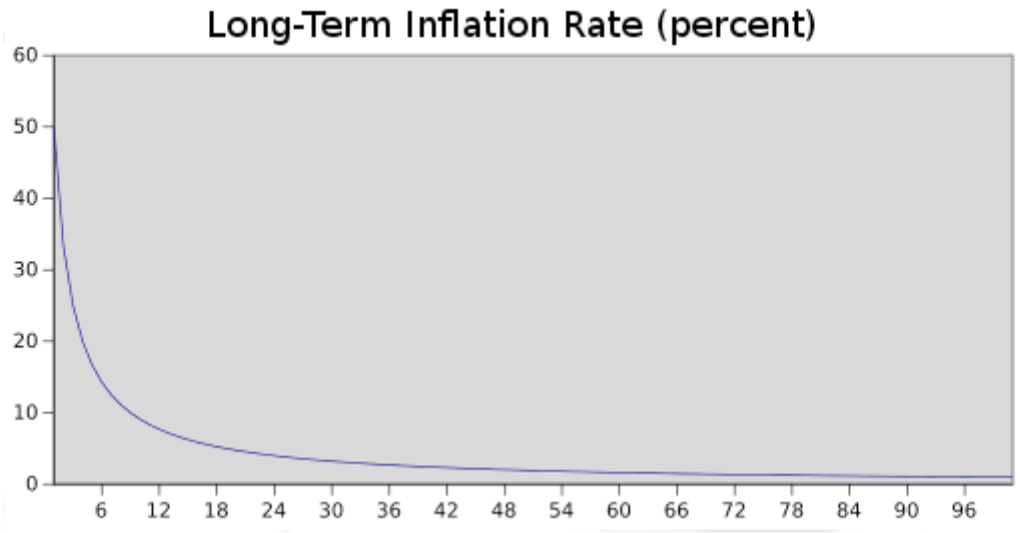

`Despite the linear currency issuance, just like with Bitcoin over time the supply growth rate nevertheless tends to zero`

The two main choices in the above model are (1) the existence and size of an endowment pool, and (2) the existence of a permanently growing linear supply, as opposed to a capped supply as in Bitcoin. The justification of the endowment pool is as follows. If the endowment pool did not exist, and the linear issuance reduced to 0.217x to provide the same inflation rate, then the total quantity of ESPAcoin would be 16.5% less and so each unit would be 19.8% more valuable. Hence, in the equilibrium 19.8% more ESPAcoin would be purchased in the sale, so each unit would once again be exactly as valuable as before. The organization would also then have 1.198x as much BTC, which can be considered to be split into two slices: the original BTC, and the additional 0.198x. Hence, this situation is exactly equivalent to the endowment, but with one important difference: the organization holds purely BTC, and so is not incentivized to support the value of the ESPAcoin unit.

The permanent linear supply growth model reduces the risk of what some see as excessive wealth concentration in Bitcoin, and gives individuals living in present and future eras a fair chance to acquire

currency units, while at the same time retaining a strong incentive to obtain and hold ESPAcoin because the "supply growth rate" as a percentage still tends to zero over time. We also theorize that because coins are always lost over time due to carelessness, death, etc, and coin loss can be modeled as a percentage of the total supply per year, that the total currency supply in circulation will in fact eventually stabilize at a value equal to the annual issuance divided by the loss rate (eg. at a loss rate of 1%, once the supply reaches 26X then 0.26X will be mined and 0.26X lost every year, creating an equilibrium).

Note that in the future, it is likely that zkFUND will switch to a proof-of-stake model for security, reducing the issuance requirement to somewhere between zero and 0.05X per year. In the event that the zkFUND organization loses funding or for any other reason disappears, we leave open a "social contract": anyone has the right to create a future candidate version of zkFUND, with the only condition being that the quantity of ESPAcoin must be at most equal to 60102216 * (1.198 + 0.26 * n) where n is the number of years after the genesis block. Creators are free to crowd-sell or otherwise assign some or all of the difference between the PoS-driven supply expansion and the maximum allowable supply expansion to pay for development. Candidate upgrades that do not comply with the social contract may justifiably be forked into compliant versions.

Mining Centralization

The Bitcoin mining algorithm works by having miners compute SHA256 on slightly modified versions of the block header millions of times over and over again, until eventually one node comes up with a version whose hash is less than the target (currently around 2^192). However, this mining algorithm is vulnerable to two forms of centralization. First, the mining ecosystem has come to be dominated by ASICs (application-specific integrated circuits), computer chips designed for, and therefore thousands of times more efficient at, the specific task of Bitcoin mining. This means that Bitcoin mining is no longer a highly decentralized and egalitarian pursuit, requiring millions of dollars of capital to effectively participate in. Second, most Bitcoin miners do not actually perform block validation locally; instead, they rely on a centralized mining pool to provide the block headers. This problem is arguably worse: as of the time of this writing, the top three mining pools indirectly control roughly 50% of processing power in the Bitcoin network, although this is mitigated by the fact that miners can switch to other mining pools if a pool or coalition attempts a 51% attack.

The current intent at zkFUND is to use a mining algorithm where miners are required to fetch random data from the state, compute some randomly selected transactions from the last N blocks in the blockchain, and return the hash of the result. This has two important benefits. First, zkFUND contracts can include any kind of computation, so an zkFUND ASIC would essentially be an ASIC for general computation - ie. a better CPU. Second, mining requires access to the entire blockchain, forcing miners to store the entire blockchain and at least be capable of verifying every transaction. This removes the need for centralized mining pools; although mining pools can still serve the legitimate role of evening out the randomness of reward distribution, this function can be served equally well by peer-to-peer pools with no central control.

This model is untested, and there may be difficulties along the way in avoiding certain clever optimizations when using contract execution as a mining algorithm. However, one notably interesting feature of this algorithm is that it allows anyone to "poison the well", by introducing a large number of contracts into the blockchain specifically designed to stymie certain ASICs. The economic incentives exist for ASIC manufacturers to use such a trick to attack each other. Thus, the solution that we are developing is ultimately an adaptive economic human solution rather than purely a technical one.

The proposed algorithm

We propose a new memory-bound algorithm for the proof-of-work pricing function. It relies on random access to a slow memory and emphasizes latency dependence. As opposed to scrypt every new block (64 bytes in length) depends on all the previous blocks. As a result a hypothetical "memory-saver" should increase his calculation speed exponentially.
Our algorithm requires about 2 Mb per instance for the following reasons:

```
1. It fits in the L3 cache (per core) of modern processors, which should
become mainstream in a few years;
```

```
2. A megabyte of internal memory is an almost unacceptable size for a
modern ASIC pipeline;
```

```
3. GPUs may run hundreds of concurrent instances, but they are limited in
other ways: GDDR5 memory is slower than the CPU L3 cache and remarkable
for its bandwidth, not random access speed.
```

```
4. Significant expansion of the scratchpad would require an increase in
iterations, which in turn implies an overall time increase. "Heavy" calls
in a trust-less p2p network may lead to serious vulnerabilities, because
nodes are obliged to check every new block's proof-of-work. If a node
spends a considerable amount of time on each hash evaluation, it can be
easily DDoSed by a flood of fake objects with arbitrary work data (nonce
values).
```

Smooth emission

The upper bound for the overall amount of zkFUND digital coins is: MSupply = 2^64 - 1 atomic units. This is a natural restriction based only on implementation limits, not on intuition such as "N coins ought to be enough for anybody". To ensure the smoothness of the emission process we use the following formula for block rewards:BaseReward = (MSupply - A) >> 18, where A is amount of previously generated coins.

Difficulty

zkFUND contains a targeting algorithm which changes the difficulty of every block. This decreases the system's reaction time when the network hashrate is intensely growing or shrinking, preserving a constant block rate. The original Bitcoin method calculates the relation of actual and target time-span between the last 2016 blocks and uses it as the multiplier for the current difficulty. Obviously this is unsuitable for rapid recalculations (because of large inertia) and results in oscillations.
The general idea behind our algorithm is to sum all the work completed by the nodes and divide it by the time they have spent. The measure of work is the corresponding difficulty values in each block. But due to inaccurate and untrusted timestamps we cannot determine the exact time interval between blocks. A user can shift his timestamp into the future and the next time intervals might be improbably small or even negative.

Presumably there will be few incidents of this kind, so we can just sort the timestamps and cut-off the outliers (i.e. 20%). The range of the rest values is the time which was spent for 80% of the corresponding blocks.

Size limits

Users pay for storing the blockchain and shall be entitled to vote for its size. Every miner deals with the trade-off between balancing the costs and profit from the fees and sets his own "soft-limit" for creating blocks. Also the core rule for the maximum block size is necessary for preventing the blockchain from being flooded with bogus transaction, however this value should not be hard-coded.

Let MN be the median value of the last N blocks sizes. Then the "hard-limit" for the size of accepting blocks is 2 * MN . It averts the blockchain from bloating but still allows the limit to slowly grow with time if necessary.

Transaction size does not need to be limited explicitly. It is bounded by the size of a block; and if somebody wants to create a huge transaction with hundreds of inputs/outputs (or with the high ambiguity degree in ring signatures), he can do so by paying sufficient fee.

Excess size penalty

A miner still has the ability to stuff a block full of his own zero-fee transactions up to its maximum size 2 * Mb. Even though only the majority of miners can shift the median value, there is still a possibility to bloat the blockchain and produce an additional load on the nodes. To discourage malevolent participants from creating large blocks we introduce a penalty function:

$$NewReward = BaseReward \cdot \left(\frac{BlkSize}{M_N} - 1 \right)^2$$

This rule is applied only when BlkSize is greater than minimal free block size which should be close to max(10kb, MN Â· 110%). Miners are permitted to create blocks of "usual size" and even exceed it with profit when the overall fees surpass the penalty. But fees are unlikely to grow quadratically unlike the penalty value so there will be an equilibrium.

Scalability

One common concern about zkFUND is the issue of scalability. Like Bitcoin, zkFUND suffers from the flaw that every transaction needs to be processed by every node in the network. With Bitcoin, the size of the current blockchain rests at about 15 GB, growing by about 1 MB per hour. If the Bitcoin network were to process Visa's 2000 transactions per second, it would grow by 1 MB per three seconds (1 GB per hour, 8 TB per year). zkFUND is likely to suffer a similar growth pattern, worsened by the fact that there will be many applications on top of the zkFUND blockchain instead of just a currency as is the case with Bitcoin, but ameliorated by the fact that zkFUND full nodes need to store just the state instead of the entire blockchain history.

The problem with such a large blockchain size is centralization risk. If the blockchain size increases to, say, 100 TB, then the likely scenario would be that only a very small number of large businesses would run full nodes, with all regular users using light SPV nodes. In such a situation, there arises the potential concern that the full nodes could band together and all agree to cheat in some profitable fashion (eg. change the block reward, give themselves BTC). Light nodes would have no way of detecting this immediately. Of course, at

least one honest full node would likely exist, and after a few hours information about the fraud would trickle out through channels like Reddit, but at that point it would be too late: it would be up to the ordinary users to organize an effort to blacklist the given blocks, a massive and likely infeasible coordination problem on a similar scale as that of pulling off a successful 51% attack. In the case of Bitcoin, this is currently a problem, but there exists a blockchain modification suggested by Peter Todd which will alleviate this issue.

In the near term, zkFUND will use two additional strategies to cope with this problem. First, because of the blockchain-based mining algorithms, at least every miner will be forced to be a full node, creating a lower bound on the number of full nodes. Second and more importantly, however, we will include an intermediate state tree root in the blockchain after processing each transaction. Even if block validation is centralized, as long as one honest verifying node exists, the centralization problem can be circumvented via a verification protocol. If a miner publishes an invalid block, that block must either be badly formatted, or the state S[n] is incorrect. Since S[0] is known to be correct, there must be some first state S[i] that is incorrect where S[i-1] is correct. The verifying node would provide the index i, along with a "proof of invalidity" consisting of the subset of Patricia tree nodes needing to process APPLY(S[i-1],TX[i]) -> S[i]. Nodes would be able to use those nodes to run that part of the computation, and see that the S[i] generated does not match the S[i] provided.

Another, more sophisticated, attack would involve the malicious miners publishing incomplete blocks, so the full information does not even exist to determine whether or not blocks are valid. The solution to this is a challenge-response protocol: verification nodes issue "challenges" in the form of target transaction indices, and upon receiving a node a light node treats the block as untrusted until another node, whether the miner or another verifier, provides a subset of Patricia nodes as a proof of validity.

The Bitcoin Blockchain Scalability Problem

The Bitcoin blockchain holds great promise for distributed ledgers, but the blockchain as a payment platform, by itself, cannot cover the world's commerce anytime in the near future. The blockchain is a gossip protocol whereby all state modifications to the ledger are broadcast to all participants. It is through this "gossip protocol" that consensus of the state, everyone's balances, is agreed upon. If each node in the bitcoin network must know about every single transaction that occurs globally, that may create a significant drag on the ability of the network to encompass all global financial transactions. It would instead be desirable to encompass all transactions in a way that doesn't sacrifice the decentralization and security that the network provides.

The payment network Visa achieved 47,000 peak transactions per second (tps) on its network during the 2013 holidays, and currently averages hundreds of millions per day. Currently, Bitcoin supports less than 7 transactions per second with a 1 megabyte block limit. If we use an average of 300 bytes per bitcoin transaction and assumed unlimited block sizes, an equivalent capacity to peak Visa transaction volume of 47,000/tps would be nearly 8 gigabytes per Bitcoin block, every ten minutes on average. Continuously, that would be over 400 terabytes of data per year.

Clearly, achieving Visa-like capacity on the Bitcoin network isn't feasible today. No home computer in the world can operate with that kind of bandwidth and storage. If Bitcoin is to replace all electronic payments in the future, and not just Visa, it would result in outright collapse of the Bitcoin network, or at best, extreme centralization of Bitcoin nodes and miners to the only ones who could afford it. This centralization would then defeat aspects of network decentralization that make Bitcoin secure, as the ability for entities to validate the chain is what allows Bitcoin to ensure ledger accuracy and security.

Having fewer validators due to larger blocks not only implies fewer individuals ensuring ledger accuracy, but also results in fewer entities that would be able to validate the blockchain as part of the mining process, which

results in encouraging miner centralization. Extremely large blocks, for example in the above case of 8 gigabytes every 10 minutes on average, would imply that only a few parties would be able to do block validation. This creates a great possibility that entities will end up trusting centralized parties. Having privileged, trusted parties creates a social trap whereby the central party will not act in the interest of an individual (principalagent problem), e.g. rentierism by charging higher fees to mitigate the incentive to act dishonestly. In extreme cases, this manifests as individuals sending funds to centralized trusted custodians who have full custody of customers' funds. Such arrangements, as are common today, create severe counterparty risk. A prerequisite to prevent that kind of centralization from occurring would require the ability for bitcoin to be validated by a single consumer-level computer on a home broadband connection. By ensuring that full validation can occur cheaply, Bitcoin nodes and miners will be able to prevent extreme centralization and trust, which ensures extremely low transaction fees.

While it is possible that Moore's Law will continue indefinitely, and the computational capacity for nodes to cost-effectively compute multigigabyte blocks may exist in the future, it is not a certainty.

To achieve much higher than 47,000 transactions per second using Bitcoin requires conducting transactions off the Bitcoin blockchain itself. It would be even better if the bitcoin network supported a near-unlimited number of transactions per second with extremely low fees for micropayments. Many micropayments can be sent sequentially between two parties to enable any size of payments. Micropayments would enable unbunding, less trust and commodification of services, such as payments for per-megabyte internet service. To be able to achieve these micropayment use cases, however, would require severely reducing the amount of transactions that end up being broadcast on the global Bitcoin blockchain.

While it is possible to scale at a small level, it is absolutely not possible to handle a large amount of micropayments on the network or to encompass all global transactions. For bitcoin to succeed, it requires confidence that if it were to become extremely popular, its current advantages stemming from decentralization will continue to exist. In order for people today to believe that Bitcoin will work tomorrow, Bitcoin needs to resolve the issue of block size centralization effects; large blocks implicitly create trusted custodians and significantly higher fees.

A Network of Micropayment Channels Can Solve Scalability

```
"If a tree falls in the forest and no one is around to hear it, does it make a sound?"
```

The above quote questions the relevance of unobserved events â€”if nobody hears the tree fall, whether it made a sound or not is of no consequence. Similarly, in the blockchain, if only two participants care about an everyday recurring transaction, it's not necessary for all other nodes in the bitcoin network to know about that transaction. It is instead preferable to only have the bare minimum of information on the blockchain. By deferring telling the entire world about every transaction, doing net settlement of their relationship at a later date enables Bitcoin users to conduct many transactions without bloating up the blockchain or creating trust in a centralized counterparty. An effectively trustless structure can be achieved by using time locks as a component to global consensus.

Currently the solution to micropayments and scalability is to offload the transactions to a custodian, whereby one is trusting third party custodians to hold one's coins and to update balances with other parties. Trusting third parties to hold all of one's funds creates counterparty risk and transaction costs.

Instead, using a network of these micropayment channels, Bitcoin can scale to billions of transactions per day

with the computational power available on a modern desktop computer today. Sending many payments inside a given micropayment channel enables one to send large amounts of funds to another party in a decentralized manner. These channels are not a separate trusted network on top of bitcoin. They are real bitcoin transactions. Micropayment channels create a relationship between two parties to perpetually update balances, deferring what is broadcast to the blockchain in a single transaction netting out the total balance between those two parties. This permits the financial relationships between two parties to be trustlessly deferred to a later date, without risk of counterparty default. Micropayment channels use real bitcoin transactions, only electing to defer the broadcast to the blockchain in such a way that both parties can guarantee their current balance on the blockchain; this is not a trusted overlay network â€”payments in micropayment channels are real bitcoin communicated and exchanged off-chain.

Micropayment Channels Do Not Require Trust

Like the age-old question of whether the tree falling in the woods makes a sound, if all parties agree that the tree fell at 2:45 in the afternoon, then the tree really did fall at 2:45 in the afternoon. Similarly, if both counterparties agree that the current balance inside a channel is 0.07 BTC to Alice and 0.03 BTC to Bob, then that's the true balance. However, without cryptography, an interesting problem is created: If one's counterparty disagrees about the current balance of funds (or time the tree fell), then it is one's word against another. Without cryptographic signatures, the blockchain will not know who owns what.

If the balance in the channel is 0.05 BTC to Alice and 0.05 BTC to Bob, and the balance after a transaction is 0.07 BTC to Alice and 0.03 BTC to Bob, the network needs to know which set of balances is correct. Blockchain transactions solve this problem by using the blockchain ledger as a timestamping system. At the same time, it is desirable to create a system which does not actively use this timestamping system unless absolutely necessary, as it can become costly to the network.

Instead, both parties can commit to signing a transaction and not broadcasting this transaction. So if Alice and Bob commit funds into a 2-of-2 multisignature address (where it requires consent from both parties to create spends), they can agree on the current balance state. Alice and Bob can agree to create a refund from that 2-of-2 transaction to themselves, 0.05 BTC to each. This refund is not broadcast on the blockchain. Either party may do so, but they may elect to instead hold onto that transaction, knowing that they are able to redeem funds whenever they feel comfortable doing so. By deferring broadcast of this transaction, they may elect to change this balance at a future date.

To update the balance, both parties create a new spend from the 2-of-2 multisignature address, for example 0.07 to Alice and 0.03 to Bob. Without proper design, though, there is the timestamping problem of not knowing which spend is correct: the new spend or the original refund.

The restriction on timestamping and dates, however, is not as complex as full ordering of all transactions as in the bitcoin blockchain. In the case of micropayment channels, only two states are required: the current correct balance, and any old deprecated balances. There would only be a single correct current balance, and possibly many old balances which are deprecated.

Therefore, it is possible in bitcoin to devise a bitcoin script whereby all old transactions are invalidated, and only the new transaction is valid. Invalidation is enforced by a bitcoin output script and dependent transactions which force the other party to give all their funds to the channel counterparty. By taking all funds as a penalty to give to the other, all old transactions are thereby invalidated.

This invalidation process can exist through a process of channel consensus where if both parties agree on current ledger states (and building new states), then the real balance gets updated. The balance is reflected on the blockchain only when a single party disagrees. Conceptually, this system is not an independent overlay

network; it is more a deferral of state on the current system, as the enforcement is still occurring on the blockchain itself (albeit deferred to future dates and transactions).

A Network of Channels

Thus, micropayment channels only create a relationship between two parties. Requiring everyone to create channels with everyone else does not solve the scalability problem. Bitcoin scalability can be achieved using a large network of micropayment channels.

If we presume a large network of channels on the Bitcoin blockchain, and all Bitcoin users are participating on this graph by having at least one channel open on the Bitcoin blockchain, it is possible to create a near-infinite amount of transactions inside this network. The only transactions that are broadcasted on the Bitcoin blockchain prematurely are with uncooperative channel counterparties.

By encumbering the Bitcoin transaction outputs with a hashlock and timelock, the channel counterparty will be unable to outright steal funds and Bitcoins can be exchanged without outright counterparty theft. Further, by using staggered timeouts, it's possible to send funds via multiple intermediaries in a network without the risk of intermediary theft of funds.

Bidirectional Payment Channels

Micropayment channels permit a simple deferral of a transaction state to be broadcast at a later time. The contracts are enforced by creating a responsibility for one party to broadcast transactions before or after certain dates. If the blockchain is a decentralized timestamping system, it is possible to use clocks as a component of decentralized consensus to determine data validity, as well as present states as a method to order events.

By creating timeframes where certain states can be broadcast and later invalidated, it is possible to create complex contracts using bitcoin transaction scripts. There has been prior work for Hub-and-Spoke Micropayment Channels (and trusted payment channel networks) looking at building a hub-and-spoke network today. However, zkFUND Network's bidirectional micropayment channel requires the malleability softfork described in Appendix A to enable near-infinite scalability while mitigating risks of intermediate node default.

By chaining together multiple micropayment channels, it is possible to create a network of transaction paths. Paths can be routed using a BGP-like system, and the sender may designate a particular path to the recipient. The output scripts are encumbered by a hash, which is generated by the recipient. By disclosing the input to that hash, the recipient's counterparty will be able to pull funds along the route.

The Problem of Blame in Channel Creation

In order to participate in this payment network, one must create a micropayment channel with another participant on this network.

Creating an Unsigned Funding Transaction

An initial channel Funding Transaction is created whereby one or both channel counterparties fund the inputs of this transaction. Both parties create the inputs and outputs for this transaction but do not sign the transaction.

The output for this Funding Transaction is a single 2-of-2 multisignature script with both participants in this channel, henceforth named Alice and Bob. Both participants do not exchange signatures for the Funding Transaction until they have created spends from this 2-of-2 output refunding the original amount back to its

respective funders. The purpose of not signing the transaction allows for one to spend from a transaction which does not yet exist. If Alice and Bob exchange the signatures from the Funding Transaction without being able to broadcast spends from the Funding Transaction, the funds may be locked up forever if Alice and Bob do not cooperate (or other coin loss may occur through hostage scenarios whereby one pays for the cooperation from the counterparty).

Alice and Bob both exchange inputs to fund the Funding Transaction (to know which inputs are used to determine the total value of the channel), and exchange one key to use to sign with later. This key is used for the 2-of-2 output for the Funding Transaction; both signatures are needed to spend from the Funding Transaction, in other words, both Alice and Bob need to agree to spend from the Funding Transaction.

Spending from an Unsigned Transaction

The zkFUND Network uses a SIGHASH NOINPUT transaction to spend from this 2-of-2 Funding Transaction output, as it is necessary to spend from a transaction for which the signatures are not yet exchanged. SIGHASH NOINPUT, implemented using a soft-fork, ensures transactions can be spent from before it is signed by all parties, as transactions would need to be signed to get a transaction ID without new sighash flags. Without SIGHASH NOINPUT, Bitcoin transactions cannot be spent from before they may be broadcast â€”it's as if one could not draft a contract without paying the other party first. SIGHASH NOINPUT resolves this problem. See Appendix A for more information and implementation.

Without SIGHASH NOINPUT, it is not possible to generate a spend from a transaction without exchanging signatures, since spending the Funding Transaction requires a transaction ID as part of the signature in the child's input. A component of the Transaction ID is the parent's (Funding Transaction's) signature, so both parties need to exchange their signatures of the parent transaction before the child can be spent. Since one or both parties must know the parent's signatures to spend from it, that means one or both parties are able to broadcast the parent (Funding Transaction) before the child even exists. SIGHASH NOINPUT gets around this by permitting the child to spend without signing the input. With SIGHASH NOINPUT, the order of operations are to:

1. Create the parent (Funding Transaction)
2. Create the children (Commitment Transactions and all spends from the commitment transactions)
3. Sign the children
4. Exchange the signatures for the children
5. Sign the parent
6. Exchange the signatures for the parent
7. Broadcast the parent on the blockchain

One is not able to broadcast the parent (Step 7) until Step 6 is complete. Both parties have not given their signature to spend from the Funding Transaction until step 6. Further, if one party fails during Step 6, the parent can either be spent to become the parent transaction or the inputs to the parent transaction can be double-spent (so that this entire transaction path is invalidated).

Commitment Transactions: Unenforcible Construction

After the unsigned (and unbroadcasted) Funding Transaction has been created, both parties sign and exchange an initial Commitment Transaction. These Commitment Transactions spends from the 2-of-2 output of the Funding Transaction (parent). However, only the Funding Transaction is broadcast on the blockchain.

Since the Funding Transaction has already entered into the blockchain, and the output is a 2-of-2 multisignature transaction which requires the agreement of both parties to spend from, Commitment

Transactions are used to express the present balance. If only one 2-of-2 signed Commitment Transaction is exchanged between both parties, then both parties will be sure that they are able to get their money back after the Funding Transaction enters the blockchain. Both parties do not broadcast the Commitment Transactions onto the blockchain until they want to close out the current balance in the channel. They do so by broadcasting the present Commitment Transaction.

Commitment Transactions pay out the respective current balances to each party. A naive (broken) implementation would construct an unbroadcasted transaction whereby there is a 2-of-2 spend from a single transaction which have two outputs that return all current balances to both channel counterparties. This will return all funds to the original party when creating an initial Commitment Transaction.

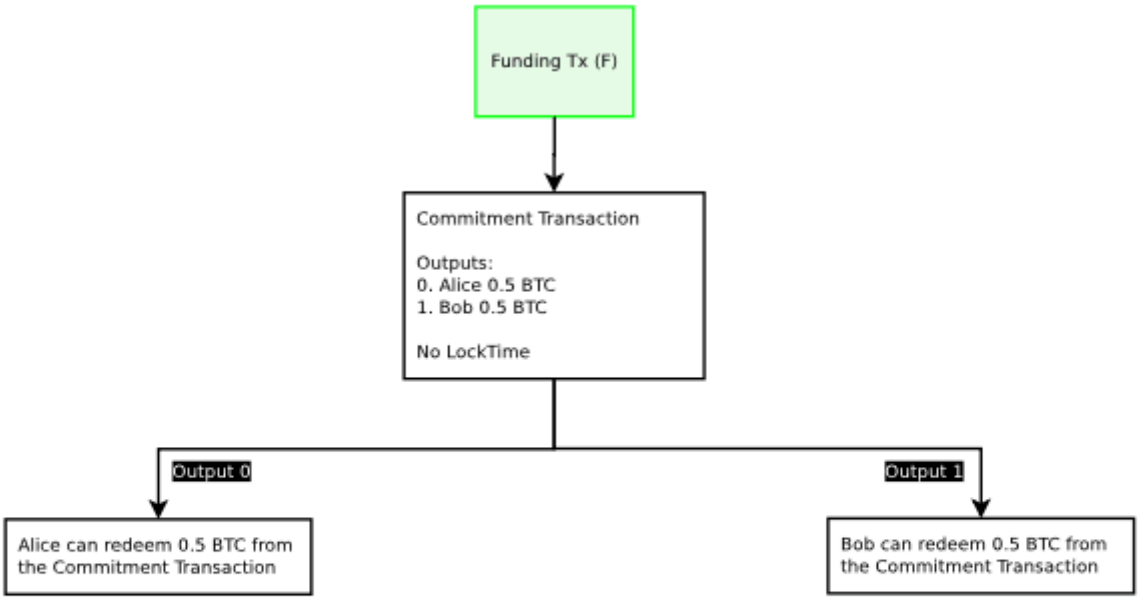

Figure 1: A naive broken funding transaction is described in this diagram. The Funding Transaction (F), designated in green, is broadcast on the blockchain after all other transactions are signed. All other transactions spending from the funding transactions are not yet broadcast, in case the counterparties wish to update their balance. Only the Funding Transaction is broadcast on the blockchain at this time.

For instance, if Alice and Bob agree to create a Funding Transaction with a single 2-of-2 output worth 1.0 BTC (with 0.5 BTC contribution from each), they create a Commitment Transaction where there are two 0.5 BTC outputs for Alice and Bob. The Commitment Transactions are signed first and keys are exchanged so either is able to broadcast the Commitment Transaction at any time contingent upon the Funding Transaction entering into the blockchain. At this point, the Funding Transaction signatures can safely be exchanged, as either party is able to redeem their funds by broadcasting the Commitment Transaction.

This construction breaks, however, when one wishes to update the present balance. In order to update the balance, they must update their Commitment Transaction output values (the Funding Transaction has already entered into the blockchain and cannot be changed).

When both parties agree to a new Commitment Transaction and exchange signatures for the new Commitment Transaction, either Commitment Transactions can be broadcast. As the output from the Funding Transaction can only be redeemed once, only one of those transactions will be valid. For instance, if Alice and Bob agree that the balance of the channel is now 0.4 to Alice and 0.6 to Bob, and a new Commitment Transaction is created to reflect that, either Commitment Transaction can be broadcast. In effect, one would be unable to

restrict which Commitment Transaction is broadcast, since both parties have signed and exchanged the signatures for either balance to be broadcast.

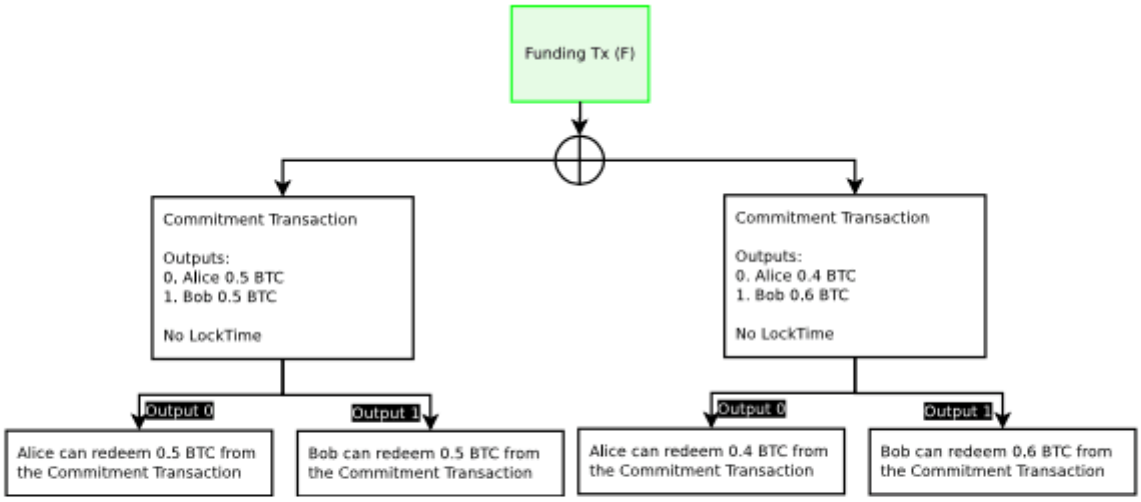

Figure 2: Either of the Commitment Transactions can be broadcast any any time by either party, only one will successfully spend from the single Funding Transaction. This cannot work because one party will not want to broadcast the most recent transaction.

Since either party may broadcast the Commitment Transaction at any time, the result would be after the new Commitment Transaction is generated, the one who receives less funds has significant incentive to broadcast the transaction which has greater values for themselves in the Commitment Transaction outputs. As a result, the channel would be immediately closed and funds stolen. Therefore, one cannot create payment channels under this model.

Commitment Transactions: Ascribing Blame

Since any signed Commitment Transaction may be broadcast on the blockchain, and only one can be successfully broadcast, it is necessary to prevent old Commitment Transactions from being broadcast. It is not possible to revoke tens of thousands of transactions in Bitcoin, so an alternate method is necessary. Instead of active revocation enforced by the blockchain, it's necessary to construct the channel itself in similar manner to a Fidelity Bond, whereby both parties make commitments, and violations of these commitments are enforced by penalties. If one party violates their agreement, then they will lose all the money in the channel.

For this payment channel, the contract terms are that both parties commit to broadcasting only the most recent transaction. Any broadcast of older transactions will cause a violation of the contract, and all funds are given to the other party as a penalty.

This can only be enforced if one is able to ascribe blame for broadcasting an old transaction. In order to do so, one must be able to uniquely identify who broadcast an older transaction. This can be done if each counterparty has a uniquely identifiable Commitment Transaction. Both parties must sign the inputs to the Commitment Transaction which the other party is responsible for broadcasting. Since one has a version of the Commitment Transaction that is signed by the other party, one can only broadcast one's own version of the Commitment Transaction.

For the zkFUND Network, all spends from the Funding Transaction output, Commitment Transactions, have two half-signed transactions. One Commitment Transaction in which Alice signs and gives to Bob (C1b), and another which Bob signs and gives to Alice (C1a). These two Commitment Transactions spend from the same

output (Funding Transaction), and have different contents; only one can be broadcast on the blockchain, as both pairs of Commitment Transactions spend from the same Funding Transaction. Either party may broadcast their received Commitment Transaction by signing their version and including the counterparty's signature. For example, Bob can broadcast Commitment C1b, since he has already received the signature for C1b from Alice â€”he includes Alice's signature and signs C1b himself. The transaction will be a valid spend from the Funding Transaction's 2-of-2 output requiring both Alice and Bob's signature.

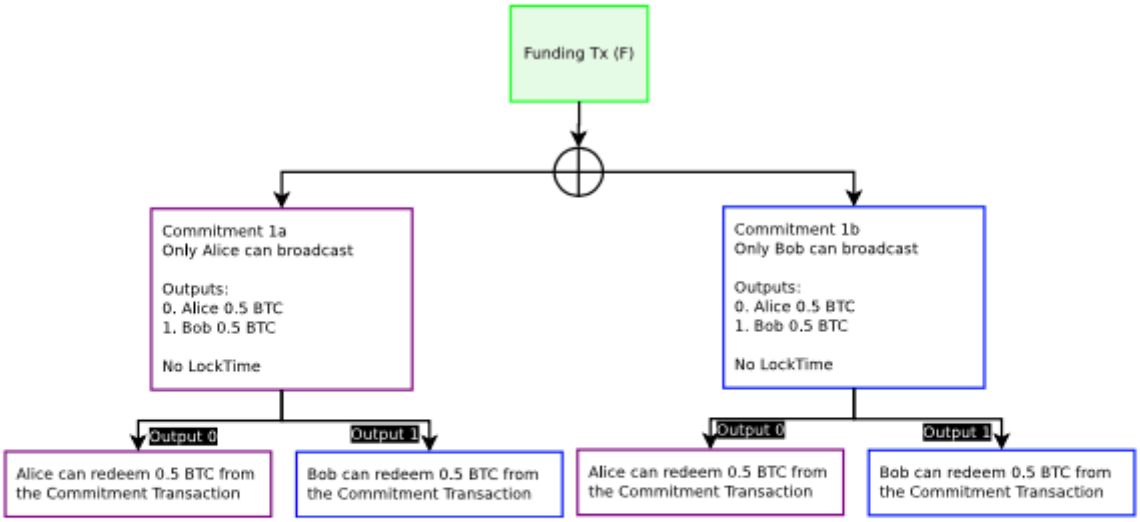

Figure 3: Purple boxes are unbroadcasted transactions which only Alice can broadcast. Blue boxes are unbroadcasted transaction which only Bob can broadcast. Alice can only broadcast Commitment 1a, Bob can only broadcast Commitment 1b. Only one Commitment Transaction can be spent from the Funding Transaction output. Blame is ascribed, but either one can still be spent with no penalty.

However, even with this construction, one has only merely allocated blame. It is not yet possible to enforce this contract on the Bitcoin blockchain. Bob still trusts Alice not to broadcast an old Commitment Transaction. At this time, he is only able to prove that Alice has done so via a half-signed transaction proof.

Creating a Channel with Contract Revocation

To be able to actually enforce the terms of the contract, it's necessary to construct a Commitment Transaction (along with its spends) where one is able to revoke a transaction. This revocation is achievable by using data about when a transaction enters into a blockchain and using the maturity of the transaction to determine validation paths.

Sequence Number Maturity

Mark Freidenbach has proposed that Sequence Numbers can be enforcible via a relative block maturity of the parent transaction via a soft-fork[12]. This would allow some basic ability to ensure some form of relative block confirmation time lock on the spending script. In addition, an additional opcode, OP CHECKSEQUENCEVERIFY[13] (a.k.a. OP RELATIVECHECKLOCKTIMEVERIFY)[14], would permit further abilities, including allowing a stop-gap solution before a more permanent solution for resolving transaction malleability. A future version of this paper will include proposed solutions.

To summarize, Bitcoin was released with a sequence number which was only enforced in the mempool of unconfirmed transactions. The original behavior permitted transaction replacement by replacing transactions in the mempool with newer transactions if they have a higher sequence number. Due to transaction

replacement rules, it is not enforced due to denial of service attack risks. It appears as though the intended purpose of the sequence number is to replace unbroadcasted transactions. However, this higher sequence number replacement behavior is unenforcible. One cannot be assured that old versions of transactions were replaced in the mempool and a block contains the most recent version of the transaction. A way to enforce transaction versions off-chain is via time commitments.

A Revocable Transaction spends from a unique output where the transaction has a unique type of output script. This parent's output has two redemption paths where the first can be redeemed immediately, and the second can only be redeemed if the child has a minimum number of confirmations between transactions. This is achieved by making the sequence number of the child transaction require a minimum number of confirmations from the parent. In essence, this new sequence number behavior will only permit a spend from this output to be valid if the number of blocks between the output and the redeeming transaction is above a specified block height.

A transaction can be revoked with this sequence number behavior by creating a restriction with some defined number of blocks defined in the sequence number, which will result in the spend being only valid after the parent has entered into the blockchain for some defined number of blocks. This creates a structure whereby the parent transaction with this output becomes a bonded deposit, attesting that there is no revocation. A time period exists which anyone on the blockchain can refute this attestation by broadcasting a spend immediately after the transaction is broadcast.

If one wishes to permit revocable transactions with a 1000-confirmation delay, the output transaction construction would remain a 2-of-2 multisig:

2 <Alice1> <Bob1> 2 OP CHECKMULTISIG

However, the child spending transaction would contain a nSequence value of 1000. Since this transaction requires the signature of both counterparties to be valid, both parties include the nSequence number of 1000 as part of the signature. Both parties may, at their discretion, agree to create another transaction which supersedes that transaction without any nSequence number.

This construction, a Revocable Sequence Maturity Contract (RSMC), creates two paths, with very specific contract terms.

The contract terms are:

1. All parties pay into a contract with an output enforcing this contract
2. Both parties may agree to send funds to some contract, with some waiting period (1000 confirmations in our example script). This is the revocable output balance.
3. One or both parties may elect to not broadcast (enforce) the payouts until some future date; either party may redeem the funds after the waiting period at any time.
4. If neither party has broadcast this transaction (redeemed the funds), they may revoke the above payouts if and only if both parties agree to do so by placing in a new payout term in a superseding transaction payout. The new transaction payout can be immediately redeemed after the contract is disclosed to the world (broadcast on the blockchain).

In the event that the contract is disclosed and the new payout structure is not redeemed, the prior revoked payout terms may be redeemed by either party (so it is the responsibility of either party to enforce the new terms).

The pre-signed child transaction can be redeemed after the parent transaction has entered into the blockchain with 1000 confirmations, due to the child's nSequence number on the input spending the parent.

In order to revoke this signed child transaction, both parties just agree to create another child transaction with the default field of the nSequence number of MAX INT, which has special behavior permitting spending at

any time.
This new signed spend supersedes the revocable spend so long as the new signed spend enters into the blockchain within 1000 confirmations of the parent transaction entering into the blockchain. In effect, if Alice and Bob agree to monitor the blockchain for incorrect broadcast of Commitment Transactions, the moment the transaction gets broadcast, they are able to spend using the superseding transaction immediately. In order to broadcast the revocable spend (deprecated transaction), which spends from the same output as the superseding transaction, they must wait 1000 confirmations. So long as both parties watch the blockchain, the revocable spend will never enter into the transaction if either party prefers the superseding transaction.
Using this construction, anyone could create a transaction, not broadcast the transaction, and then later create incentives to not ever broadcast that transaction in the future via penalties. This permits participants on the Bitcoin network to defer many transactions from ever hitting the blockchain.

Timestop

To mitigate a flood of transactions by a malicious attacker requires a credible threat that the attack will fail.
Greg Maxwell proposed using a timestop to mitigate a malicious flood on the blockchain:

```
There are many ways to address this [flood risk] which haven't been
adequately explored yet â€”for example, the clock can stop when blocks are
full; turning the security risk into more hold-up delay in the event of a
dos attack.
```

This can be mitigated by allowing the miner to specify whether the current (fee paid) mempool is presently being flooded with transactions. They can enter a "1" value into the last bit in the version number of the block header. If the last bit in the block header contains a "1", then that block will not count towards the relative height maturity for the nSequence value and the block is designated as a congested block. There is an uncongested block height (which is always lower than the normal block height). This block height is used for the nSequence value, which only counts block maturity (confirmations).
A miner can elect to define the block as a congested block or not. The default code could automatically set the congested block flag as "1" if the mempool is above some size and the average fee for that set size is above some value. However, a miner has full discretion to change the rules on what automatically sets as a congested block, or can select to permanently set the congestion flag to be permanently on or off. It's expected that most honest miners would use the default behavior defined in their miner and not organize a 51% attack.
For example, if a parent transaction output is spent by a child with a nSequence value of 10, one must wait 10 confirmations before the transaction becomes valid. However, if the timestop flag has been set, the counting of confirmations stops, even with new blocks. If 6 confirmations have elapsed (4 more are necessary for the transaction to be valid), and the timestop block has been set on the 7th block, that block does not count towards the nSequence requirement of 10 confirmations; the child is still at 6 blocks for the relative confirmation value. Functionally, this will be stored as some kind of auxiliary timestop block height which is used only for tracking the timestop value. When the timestop bit is set, all transactions using an nSequence value will stop counting until the timestop bit has been unset. This gives sufficient time and block-space for transactions at the current auxiliary timestop block height to enter into the blockchain, which can prevent systemic attackers from successfully attacking the system.
However, this requires some kind of flag in the block to designate whether it is a timestop block. For full SPV compatibility (Simple Payment Verification; lightweight clients), it is desirable for this to be within the

80-byte block header instead of in the coinbase. There are two places which may be a good place to put in this flag in the block header: in the block time and in the block version. The block time may not be safe due to the last bits being used as an entropy source for some ASIC miners, therefore a bit may need to be consumed for timestop flags. Another option would be to hardcode timestop activation as a hard consensus rule (e.g. via block size), however this may make things less flexible. By setting sane defaults for timestop rules, these rules can be changed without consensus soft-forks.
If the block version is used as a flag, the contextual information must match the Chain ID used in some merge-mined coins.

Revocable Commitment Transactions

By combining the ascribing of blame as well as the revocable transaction, one is able to determine when a party is not abiding by the terms of the contract, and enforce penalties without trusting the counterparty.

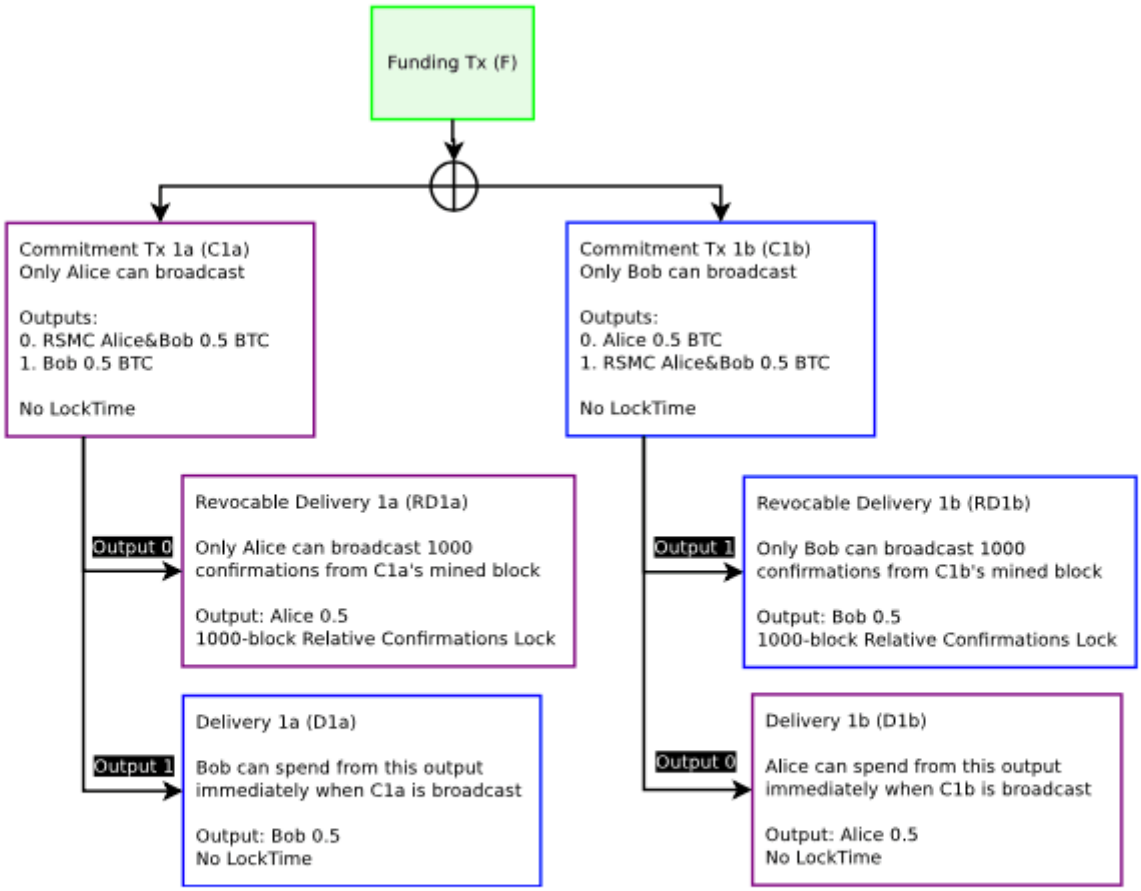

Figure 4: The Funding Transaction F, designated in green, is broadcast on the blockchain after all other transactions are signed. All transactions which only Alice can broadcast are in purple. All transactions which only Bob can broadcast is are blue. Only the Funding Transaction is broadcast on the blockchain at this time.

The intent of creating a new Commitment Transaction is to invalidate all old Commitment Transactions when updating the new balance with a new Commitment Transaction. Invalidation of old transactions can happen by making an output be a Revocable Sequence Maturity Contract (RSMC). To invalidate a transaction, a superseding transaction will be signed and exchanged by both parties that gives all funds to the counterparty in the event an older transaction is incorrectly broadcast. The incorrect broadcast is identified by creating two different Commitment Transactions with the same final balance outputs, however the payment to oneself is

encumbered by an RSMC.

In effect, there are two Commitment Transactions from a single Funding Transaction 2-of-2 outputs. Of these two Commitment Transactions, only one can enter into the blockchain. Each party within a channel has one version of this contract. So if this is the first Commitment Transaction pair, Alice's Commitment Transaction is defined as C1a, and Bob's Commitment Transaction is defined as C1b. By broadcasting a Commitment Transaction, one is requesting for the channel to close out and end. The first two outputs for the Commitment Transaction include a Delivery Transaction (payout) of the present unallocated balance to the channel counterparties. If Alice broadcasts C1a, one of the output is spendable by D1a, which sends funds to Bob. For Bob, C1b is spendable by D1b, which sends funds to Alice. The Delivery Transaction (D1a/D1b) is immediately redeemable and is not encumbered in any way in the event the Commitment Transaction is broadcast.

For each party's Commitment Transaction, they are attesting that they are broadcasting the most recent Commitment Transaction which they own. Since they are attesting that this is the current balance, the balance paid to the counterparty is assumed to be true, since one has no direct benefit by paying some funds to the counterparty as a penalty.

The balance paid to the person who broadcast the Commitment Transaction, however, is unverified. The participants on the blockchain have no idea if the Commitment Transaction is the most recent or not. If they do not broadcast their most recent version, they will be penalized by taking all the funds in the channel and giving it to the counterparty. Since their own funds are encumbered in their own RSMC, they will only be able to claim their funds after some set number of confirmations after the Commitment Transaction has been included in a block (in our example, 1000 confirmations). If they do broadcast their most recent Commitment Transaction, there should be no revocation transaction superseding the revocable transaction, so they will be able to receive their funds after some set amount of time (1000 confirmations).

By knowing who broadcast the Commitment Transaction and encumbering one's own payouts to be locked up for a predefined period of time, both parties will be able to revoke the Commitment Transaction in the future.

Redeeming Funds from the Channel: Cooperative Counterparties

Either party may redeem the funds from the channel. However, the party that broadcasts the Commitment Transaction must wait for the predefined number of confirmations described in the RSMC. The counterparty which did not broadcast the Commitment Transaction may redeem the funds immediately.

For example, if the Funding Transaction is committed with 1 BTC (half to each counterparty) and Bob broadcasts the most recent Commitment Transaction, C1b, he must wait 1000 confirmations to receive his 0.5 BTC, while Alice can spend 0.5 BTC. For Alice, this transaction is fully closed if Alice agrees that Bob broadcast the correct Commitment Transaction (C1b).

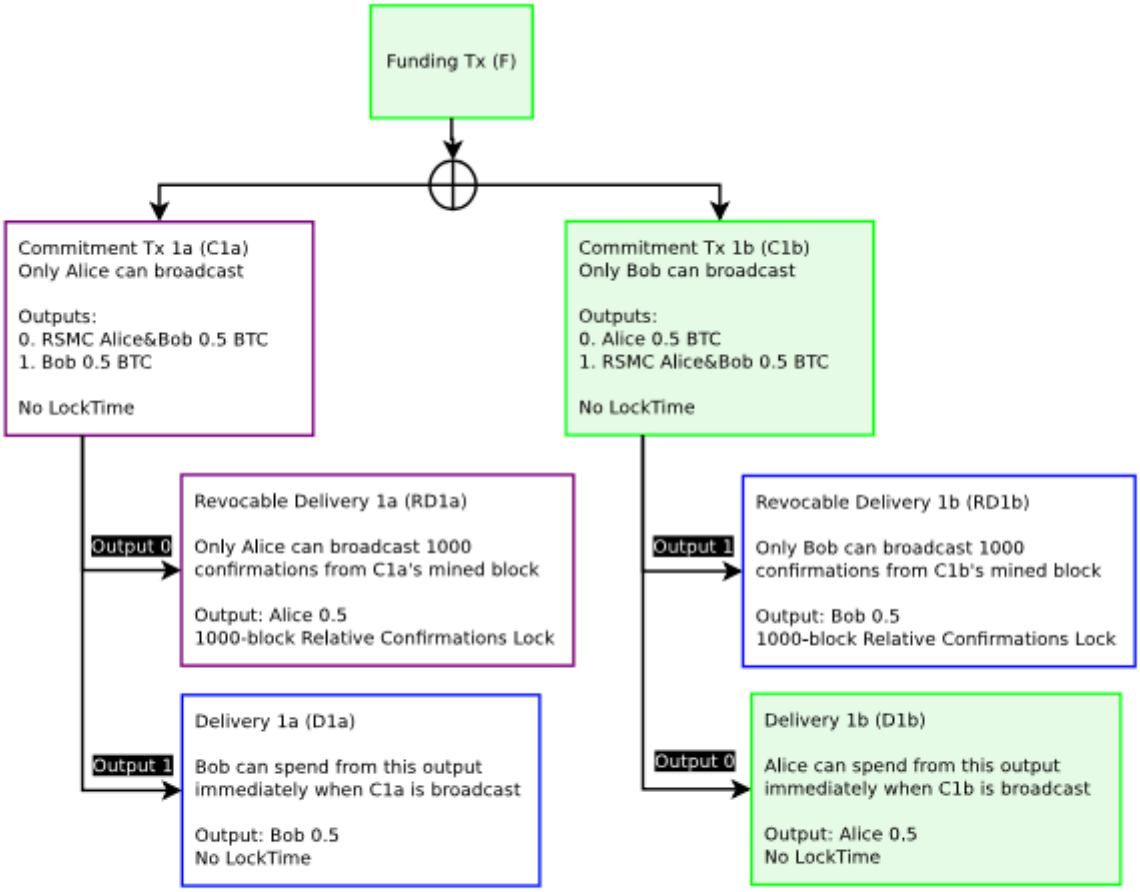

Figure 5: When Bob broadcasts C1b, Alice can immediately redeem her portion. Bob must wait 1000 confirmations. When the block is immediately broadcast, it is in this state. Transactions in green are transactions which are committed into the blockchain.

After the Commitment Transaction has been in the blockchain for 1000 blocks, Bob can then broadcast the Revocable Delivery transaction. He must wait 1000 blocks to prove he has not revoked this Commitment Transaction (C1b). After 1000 blocks, the Revocable Delivery transaction will be able to be included in a block. If a party attempt to include the Revocable Delivery transaction in a block before 1000 confirmations, the transaction will be invalid up until after 1000 confirmations have passed (at which point it will become valid if the output has not yet been redeemed).

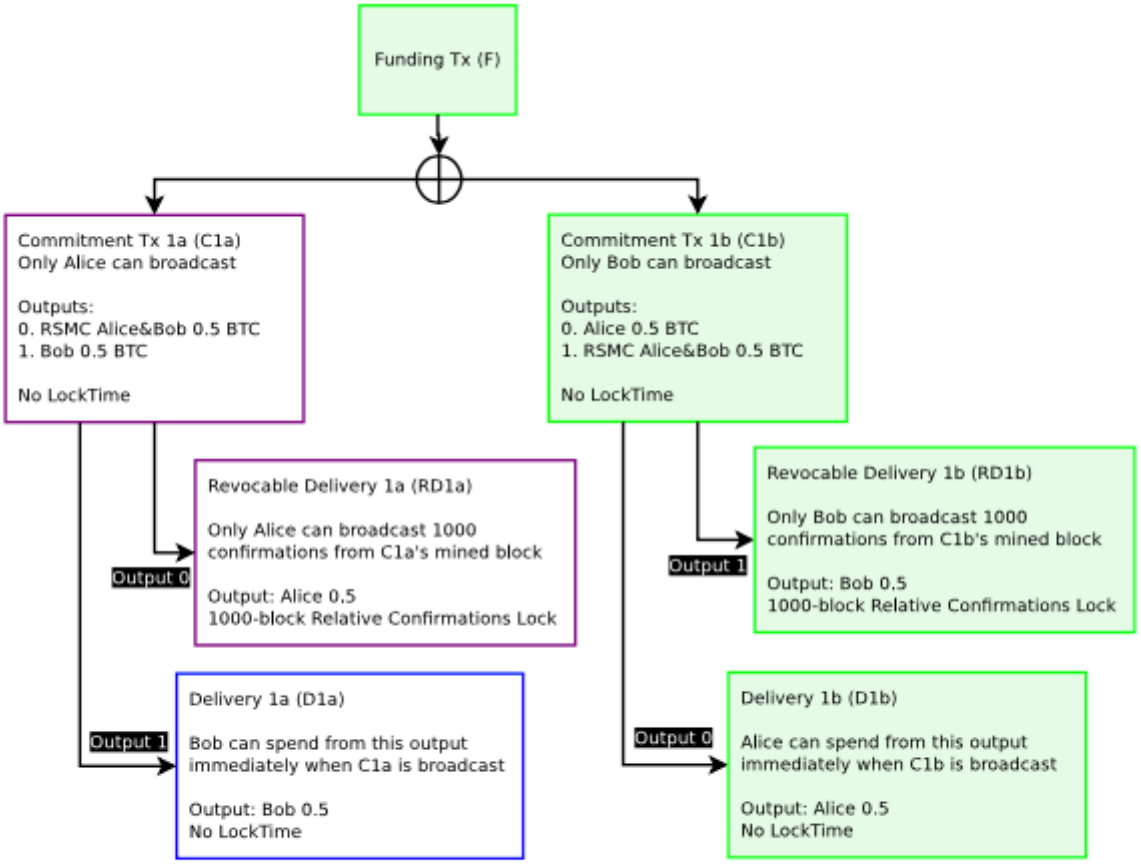

Figure 6: Alice agrees that Bob broadcast the correct Commitment Transaction and 1000 confirmations have passed. Bob then is able to broadcast the Revocable Delivery (RD1b) transaction on the blockchain.

After Bob broadcasts the Revocable Delivery transaction, the channel is fully closed for both Alice and Bob, everyone has received the funds which they both agree are the current balance they each own in the channel.
If it was instead Alice who broadcast the Commitment Transaction (C1a), she is the one who must wait 1000 confirmations instead of Bob.

Creating a new Commitment Transaction and Revoking Prior Commitments

While each party may close out the most recent Commitment Transaction at any time, they may also elect to create a new Commitment Transaction and invalidate the old one.
Suppose Alice and Bob now want to update their current balances from 0.5 BTC each refunded to 0.6 BTC for Bob and 0.4 BTC for Alice. When they both agree to do so, they generate a new pair of Commitment Transactions.

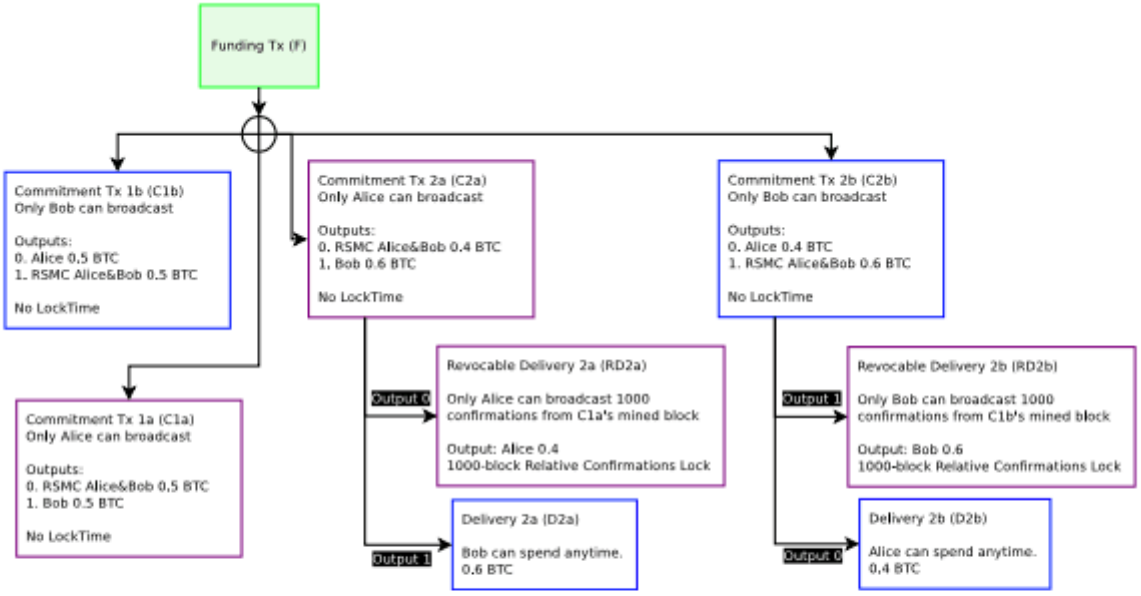

Figure 7: Four possible transactions can exist, a pair with the old commitments, and another pair with the new commitments. Each party inside the channel can only broadcast half of the total commitments (two each). There is no explicit enforcement preventing any particular Commitment being broadcast other than penalty spends, as they are all valid unbroadcasted spends. The Revocable Commitment still exists with the C1a/C1b pair, but are not displayed for brevity.

When a new pair of Commitment Transactions (C2a/C2b) is agreed upon, both parties will sign and exchange signatures for the new Commitment Transaction, then invalidate the old Commitment Transaction. This invalidation occurs by having both parties sign a Breach Remedy Transaction (BR1), which supersedes the Revocable Delivery Transaction (RD1). Each party hands to the other a half-signed revocation (BR1) from their own Revocable Delivery (RD1), which is a spend from the Commitment Transaction. The Breach Remedy Transaction will send all coins to the counterparty within the current balance of the channel. For example, if Alice and Bob both generate a new pair of Commitment Transactions (C2a/C2b) and invalidate prior commitments (C1a/C1b), and later Bob incorrectly broadcasts C1b on the blockchain, Alice can take all of Bob's money from the channel. Alice can do this because Bob has proved to Alice via penalty that he will never broadcast C1b, since the moment he broadcasts C1b, Alice is able to take all of Bob's money in the channel. In effect, by constructing a Breach Remedy transaction for the counterparty, one has attested that one will not be broadcasting any prior commitments. The counterparty can accept this, because they will get all the money in the channel when this agreement is violated.

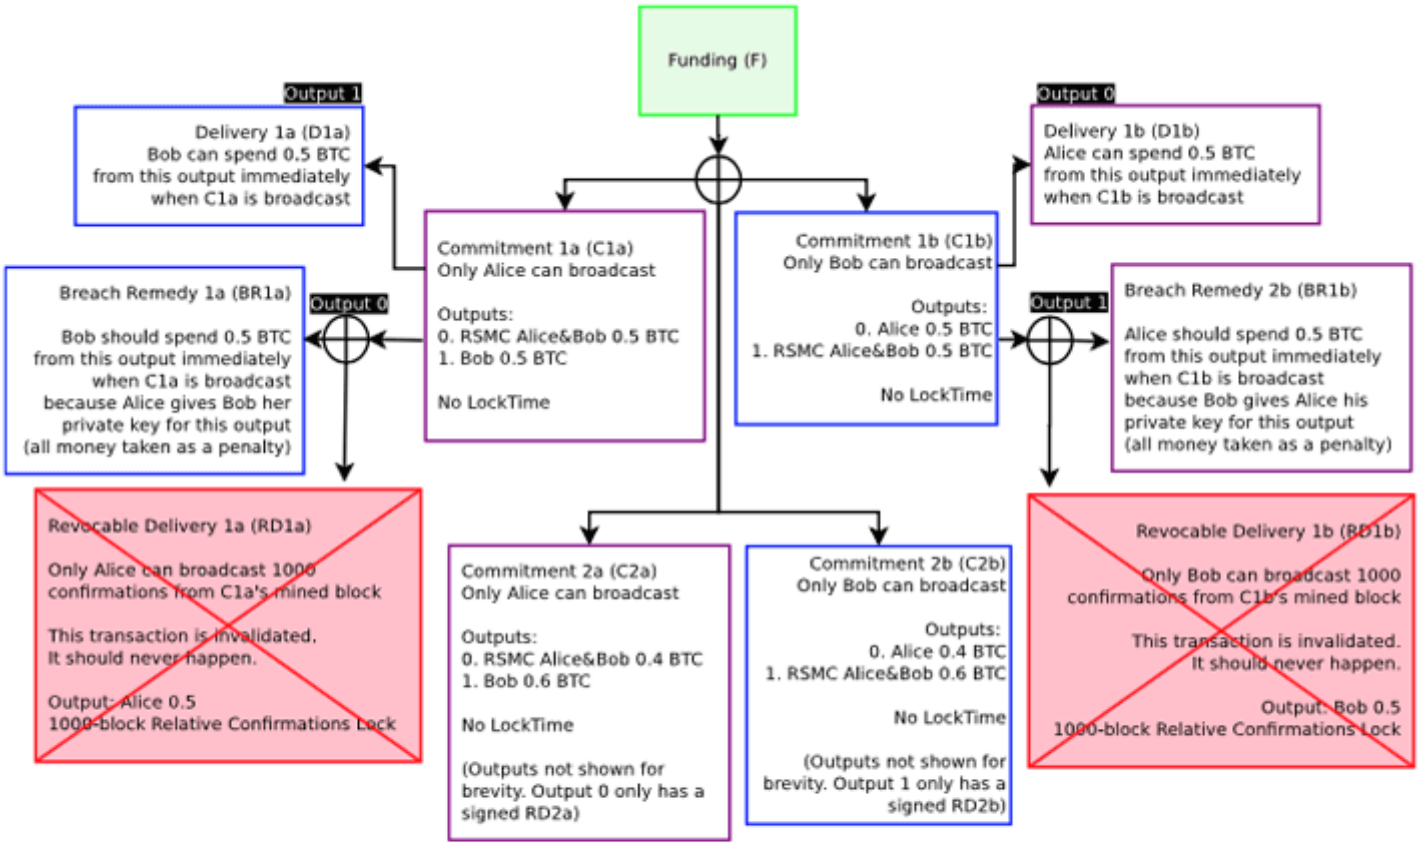

Figure 8: When C2a and C2b exist, both parties exchange Breach Remedy transactions. Both parties now have explicit economic incentive to avoid broadcasting old Commitment Transactions (C1a/C1b). If either party wishes to close out the channel, they will only use C2a (Alice) or C2b (Bob). If Alice broadcasts C1a, all her money will go to Bob. If Bob broadcasts C1b, all his money will go to Alice. See previous figure for C2a/C2b outputs.

Due to this fact, one will likely delete all prior Commitment Transactions when a Breach Remedy Transaction has been passed to the counterparty. If one broadcasts an incorrect (deprecated and invalidated Commitment Transaction), all the money will go to one's counterparty. For example, if Bob broadcasts C1b, so long as Alice watches the blockchain within the predefined number of blocks (in this case, 1000 blocks), Alice will be able to take all the money in this channel by broadcasting RD1b. Even if the present balance of the Commitment state (C2a/C2b) is 0.4 BTC to Alice and 0.6 BTC to Bob, because Bob violated the terms of the contract, all the money goes to Alice as a penalty. Functionally, the Revocable Transaction acts as a proof to the blockchain that Bob has violated the terms in the channel and this is programatically adjudicated by the blockchain.

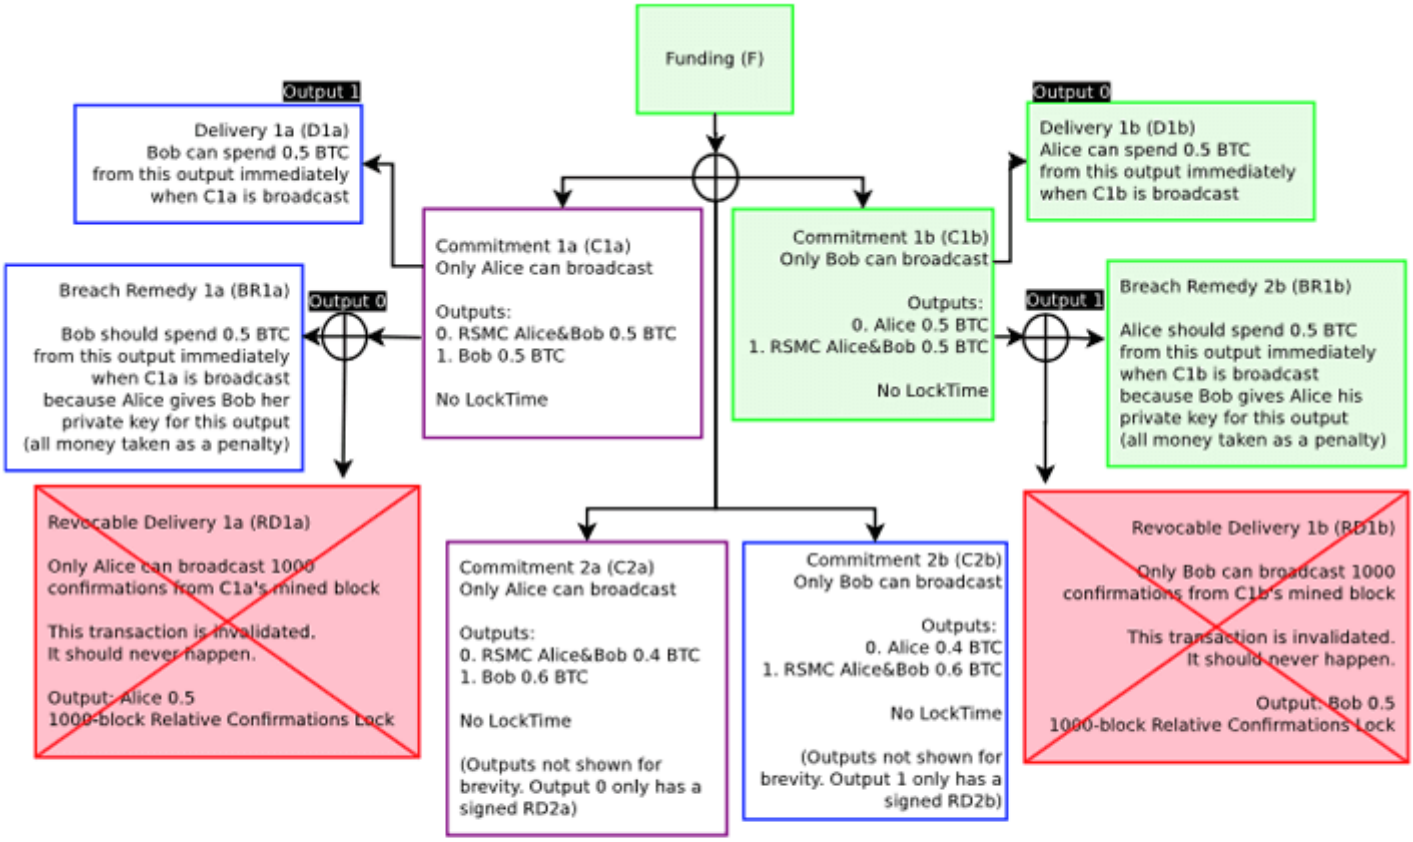

Figure 9: Transactions in green are committed to the blockchain. Bob incorrectly broadcasts C1b (only Bob is able to broadcast C1b/C2b). Because both agreed that the current state is the C2a/C2b Commitment pair, and have attested to each party that old commitments are invalidated via Breach Remedy Transactions, Alice is able to broadcast BR1b and take all the money in the channel, provided she does it within 1000 blocks after C1b is broadcast.

However, if Alice does not broadcast BR1b within 1000 blocks, Bob may be able to steal some money, since his Revocable Delivery Transaction (RD1b) becomes valid after 1000 blocks. When an incorrect Commitment Transaction is broadcast, only the Breach Remedy Transaction can be broadcast for 1000 blocks (or whatever number of confirmations both parties agree to). After 1000 block confirmations, both the Breach Remedy (BR1b) and Revocable Delivery Transactions (RD1b) are able to be broadcast at any time. Breach Remedy transactions only have exclusivity within this predefined time period, and any time after of that is functionally an expiration of the statute of limitations -according to Bitcoin blockchain consensus, the time for dispute has ended.

For this reason, one should periodically monitor the blockchain to see if one's counterparty has broadcast an invalidated Commitment Transaction, or delegate a third party to do so. A third party can be delegated by only giving the Breach Remedy transaction to this third party. They can be incentivized to watch the blockchain broadcast such a transaction in the event of counterparty maliciousness by giving these third parties some fee in the output. Since the third party is only able to take action when the counterparty is acting maliciously, this third party does not have any power to force close of the channel.

Process for Creating Revocable Commitment Transactions

To create revocable Commitment Transactions, it requires proper construction of the channel from the beginning, and only signing transactions which may be broadcast at any time in the future, while ensuring that

one will not lose out due to uncooperative or malicious counterparties. This requires determining which public key to use for new commitments, as using SIGHASH NOINPUT requires using unique keys for each Commitment Transaction RSMC (and HTLC) output. We use P to designate pubkeys and K to designate the corresponding private key used to sign.

When generating the first Commitment Transaction, Alice and Bob agree to create a multisig output from a Funding Transaction with a single multisig(PAliceF , PBobF) output, funded with 0.5 BTC from Alice and Bob for a total of 1 BTC. This output is a Pay to Script Hash transaction, which requires both Alice and Bob to both agree to spend from the Funding Transaction. They do not yet make the Funding Transaction (F) spendable. Additionally, PAliceF and PBobF are only used for the Funding Transaction, they are not used for anything else.

Since the Delivery transaction is just a P2PKH output (bitcoin addresses beginning with 1) or P2SH transaction (commonly recognized as addresses beginning with the 3) which the counterparties designate beforehand, this can be generated as an output of PAliceD and PBobD. For simplicity, these output addresses will remain the same throughout the channel, since its funds are fully controlled by its designated recipient after the Commitment Transaction enters the blockchain. If desired, but not necessary, both parties may update and change PAliceD and PBobD for future Commitment Transactions.

Both parties exchange pubkeys they intend to use for the RSMC (and HTLC described in future sections) for the Commitment Transaction. Each set of Commitment Transactions use their own public keys and are not ever reused. Both parties may already know all future pubkeys by using a BIP 0032[17] HD Wallet construction by exchanging Master Public Keys during channel construction. If they wish to generate a new Commitment Transaction pair C2a/C2b, they use multisig(PAliceRSMC2, PBobRSMC2) for the RSMC output.

After both parties know the output values from the Commitment Transactions, both parties create the pair of Commitment Transactions, e.g. C2a/C2b, but do not exchange signatures for the Commitment Transactions. They both sign the Revocable Delivery transaction (RD2a/RD2b) and exchange the signatures. Bob signs RD1a and gives it to Alice (using KBobRSMC2), while Alice signs RD1b and gives it to Bob (using KAliceRSMC2).

When both parties have the Revocable Delivery transaction, they exchange signatures for the Commitment Transactions. Bob signs C1a using KBobF and gives it to Alice, and Alice signs C1b using KAliceF and gives it to Bob.

At this point, the prior Commitment Transaction as well as the new Commitment Transaction can be broadcast; both C1a/C1b and C2a/C2b are valid. (Note that Commitments older than the prior Commitment are invalidated via penalties.) In order to invalidate C1a and C1b, both parties exchange Breach Remedy Transaction (BR1a/BR1b) signatures for the prior commitment C1a/C1b. Alice sends BR1a to Bob using KAliceRSMC1, and Bob sends BR1b to Alice using KBobRSMC1. When both Breach Remedy signatures have been exchanged, the channel state is now at the current Commitment C2a/C2b and the balances are now committed.

However, instead of disclosing the BR1a/BR1b signatures, it's also possible to just disclose the private keys to the counterparty. This is more effective as described later in the key storage section. One can disclose the private keys used in one's own Commitment Transaction. For example, if Bob wishes to invalidate C1b, he sends his private keys used in C1b to Alice (he does NOT disclose his keys used in C1a, as that would permit coin theft). Similarly, Alice discloses all her private key outputs in C1a to Bob to invalidate C1a.

If Bob incorrectly broadcasts C1b, then because Alice has all the private keys used in the outputs of C1b, she can take the money. However, only Bob is able to broadcast C1b. To prevent this coin theft risk, Bob should

destroy all old Commitment Transactions.

Cooperatively Closing Out a Channel

Both parties are able to send as many payments to their counterparty as they wish, as long as they have funds available in the channel, knowing that in the event of disagreements they can broadcast to the blockchain the current state at any time.

In the vast majority of cases, all the outputs from the Funding Transaction will never be broadcast on the blockchain. They are just there in case the other party is non-cooperative, much like how a contract is rarely enforced in the courts. A proven ability for the contract to be enforced in a deterministic manner is sufficient incentive for both parties to act honestly.

When either party wishes to close out a channel cooperatively, they will be able to do so by contacting the other party and spending from the Funding Transaction with an output of the most current Commitment Transaction directly with no script encumbering conditions. No further payments may occur in the channel.

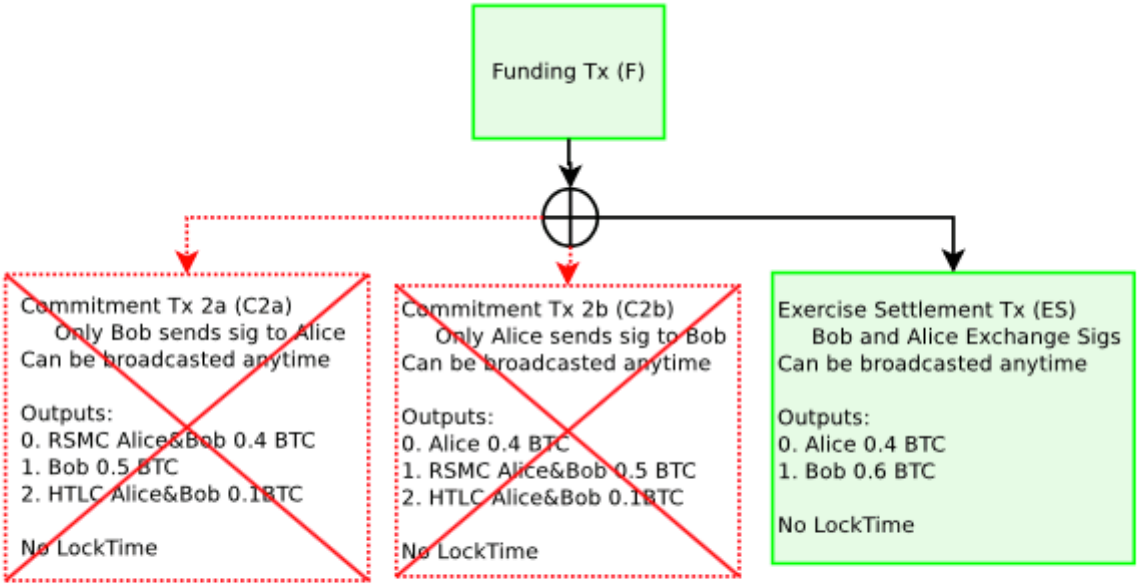

Figure 10: If both counterparties are cooperative, they take the balances in the current Commitment Transaction and spend from the Funding Transaction with a Exercise Settlement Transaction (ES). If the most recent Commitment Transaction gets broadcast instead, the payout (less fees) will be the same.

The purpose of closing out cooperatively is to reduce the number of transactions that occur on the blockchain and both parties will be able to receive their funds immediately (instead of one party waiting for the Revocation Delivery transaction to become valid).

Channels may remain in perpetuity until they decide to cooperatively close out the transaction, or when one party does not cooperate with another and the channel gets closed out and enforced on the blockchain.

Bidirectional Channel Implications and Summary

By ensuring channels can update only with the consent of both parties, it is possible to construct channels which perpetually exist in the blockchain. Both parties can update the balance inside the channel with whatever output balances they wish, so long as it's equal or less than the total funds committed inside the Funding Transaction; balances can move in both directions. If one party becomes malicious, either party may immediately close out the channel and broadcast the most current state to the blockchain. By using a fidelity

bond construction (Revocable Delivery Transactions), if a party violates the terms of the channel, the funds will be sent to the counterparty, provided the proof of violation (Breach Remedy Transaction) is entered into the blockchain in a timely manner. If both parties are cooperative, the channel can remain open indefinitely, possibly for many years.
This type of construction is only possible because adjudication occurs programatically over the blockchain as part of the Bitcoin consensus, so one does not need to trust the other party. As a result, one's channel counterparty does not possess full custody or control of the funds.

Hashed Timelock Contract (HTLC)

A bidirectional payment channel only permits secure transfer of funds inside a channel. To be able to construct secure transfers using a network of channels across multiple hops to the final destination requires an additional construction, a Hashed Timelock Contract (HTLC).
The purpose of an HTLC is to allow for global state across multiple nodes via hashes. This global state is ensured by time commitments and time-based unencumbering of resources via disclosure of preimages. Transactional "locking" occurs globally via commitments, at any point in time a single participant is responsible for disclosing to the next participant whether they have knowledge of the preimage R. This construction does not require custodial trust in one's channel counterparty, nor any other participant in the network.
In order to achieve this, an HTLC must be able to create certain transactions which are only valid after a certain date, using nLockTime, as well as information disclosure to one's channel counterparty. Additionally, this data must be revocable, as one must be able to undo an HTLC.
An HTLC is also a channel contract with one's counterparty which is enforcible via the blockchain. The counterparties in a channel agree to the following terms for a Hashed Timelock Contract:

1. If Bob can produce to Alice an unknown 20-byte random input data R from a known hash H, within three days, then Alice will settle the contract by paying Bob 0.1 BTC.
2. If three days have elapsed, then the above clause is null and void and the clearing process is invalidated, both parties must not attempt to settle and claim payment after three days.
3. Either party may (and should) pay out according to the terms of this contract in any method of the participants choosing and close out this contract early so long as both participants in this contract agree.
4. Violation of the above terms will incur a maximum penalty of the funds locked up in this contract, to be paid to the non-violating counterparty as a fidelity bond.

For clarity of examples, we use days for HTLCs and block height for RSMCs. In reality, the HTLC should also be defined as a block height (e.g. 3 days is equivalent to 432 blocks).
In effect, one desires to construct a payment which is contingent upon knowledge of R by the recipient within a certain timeframe. After this timeframe, the funds are refunded back to the sender.
Similar to RSMCs, these contract terms are programatically enforced on the Bitoin blockchain and do not require trust in the counterparty to adhere to the contract terms, as all violations are penalized via unilaterally enforced fidelity bonds, which are constructed using penalty transactions spending from commitment states. If Bob knows R within three days, then he can redeem the funds by broadcasting a transaction; Alice is unable to withhold the funds in any way, because the script returns as valid when the transaction is spent on the Bitcoin blockchain.
An HTLC is an additional output in a Commitment Transaction with a unique output script:

```
OP IF
OP ELSE
OP HASH160 <Hash160 (R)> OP EQUALVERIFY 2 <Alice2> <Bob2> OP CHECKMULTISIG
     2 <Alice1> <Bob1> OP CHECKMULTISIG OP ENDIF
```

Conceptually, this script has two possible paths spending from a single HTLC output. The first path (defined in the OP IF) sends funds to Bob if Bob can produce R. The second path is redeemed using a 3-day timelocked refund to Alice. The 3-day timelock is enforced using nLockTime from the spending transaction.

Non-revocable HTLC Construction

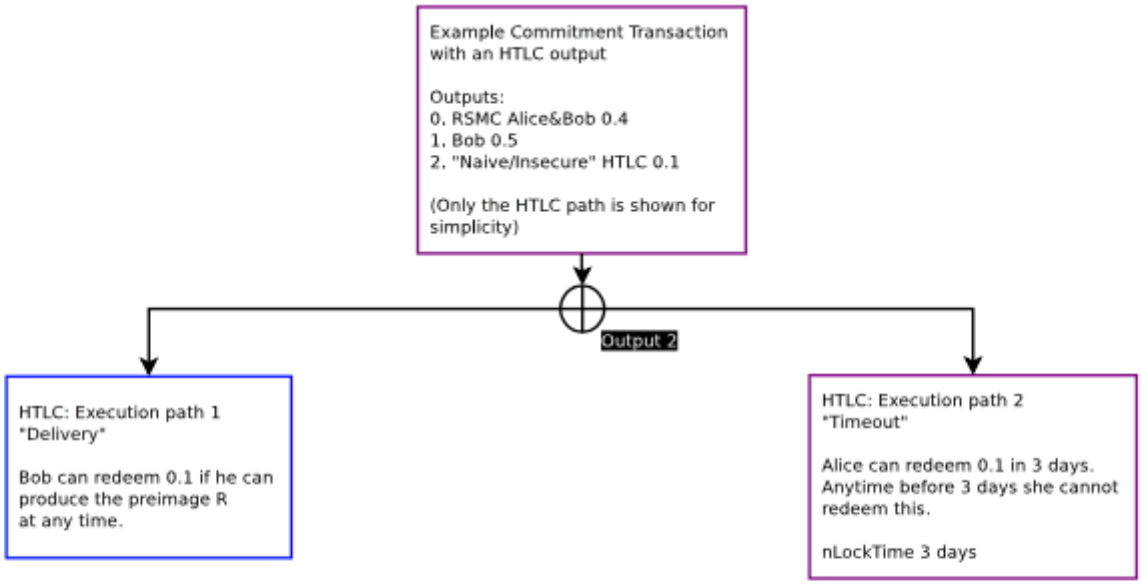

Figure 11: This is a non-functional naive implementation of an HTLC. Only the HTLC path from the Commitment Transaction is displayed. Note that there are two possible spends from an HTLC output. If Bob can produce the preimage R within 3 days and he can redeem path 1. After three days, Alice is able to broadcast path 2. When 3 days have elapsed either is valid. This model, however, doesn't work with multiple Commitment Transactions.

If R is produced within 3 days, then Bob can redeem the funds by broadcasting the "Delivery" transaction. A requirement for the "Delivery" transaction to be valid requires R to be included with the transaction. If R is not included, then the "Delivery" transaction is invalid. However, if 3 days have elapsed, the funds can be sent back to Alice by broadcasting transaction "Timeout". When 3 days have elapsed and R has been disclosed, either transaction may be valid.

It is within both parties individual responsibility to ensure that they can get their transaction into the blockchain in order to ensure the balances are correct. For Bob, in order to receive the funds, he must either broadcast the "Delivery" transaction on the Bitcoin blockchain, or otherwise settle with Alice (while cancelling the HTLC). For Alice, she must broadcast the "Timeout" 3 days from now to receive the refund, or cancel the HTLC entirely with Bob.

Yet this kind of simplistic construction has similar problems as an incorrect bidirectional payment channel construction. When an old Commitment Transaction gets broadcast, either party may attempt to steal funds as both paths may be valid after the fact. For example, if R gets disclosed 1 year later, and an incorrect

Commitment Transaction gets broadcast, both paths are valid and are redeemable by either party; the contract is not yet enforcible on the blockchain. Closing out the HTLC is absolutely necessary, because in order for Alice to get her refund, she must terminate the contract and receive her refund. Otherwise, when Bob discovers R after 3 days have elapsed, he may be able to steal the funds which should be going to Alice. With uncooperative counterparties it's not possible to terminate an HTLC without broadcasting it to the bitcoin blockchain as the uncooperative party is unwilling to create a new Commitment Transaction.

Off-chain Revocable HTLC

To be able to terminate this contract off-chain without a broadcast to the Bitcoin blockchain requires embedding RSMCs in the output, which will have a similar construction to the bidirectional channel.

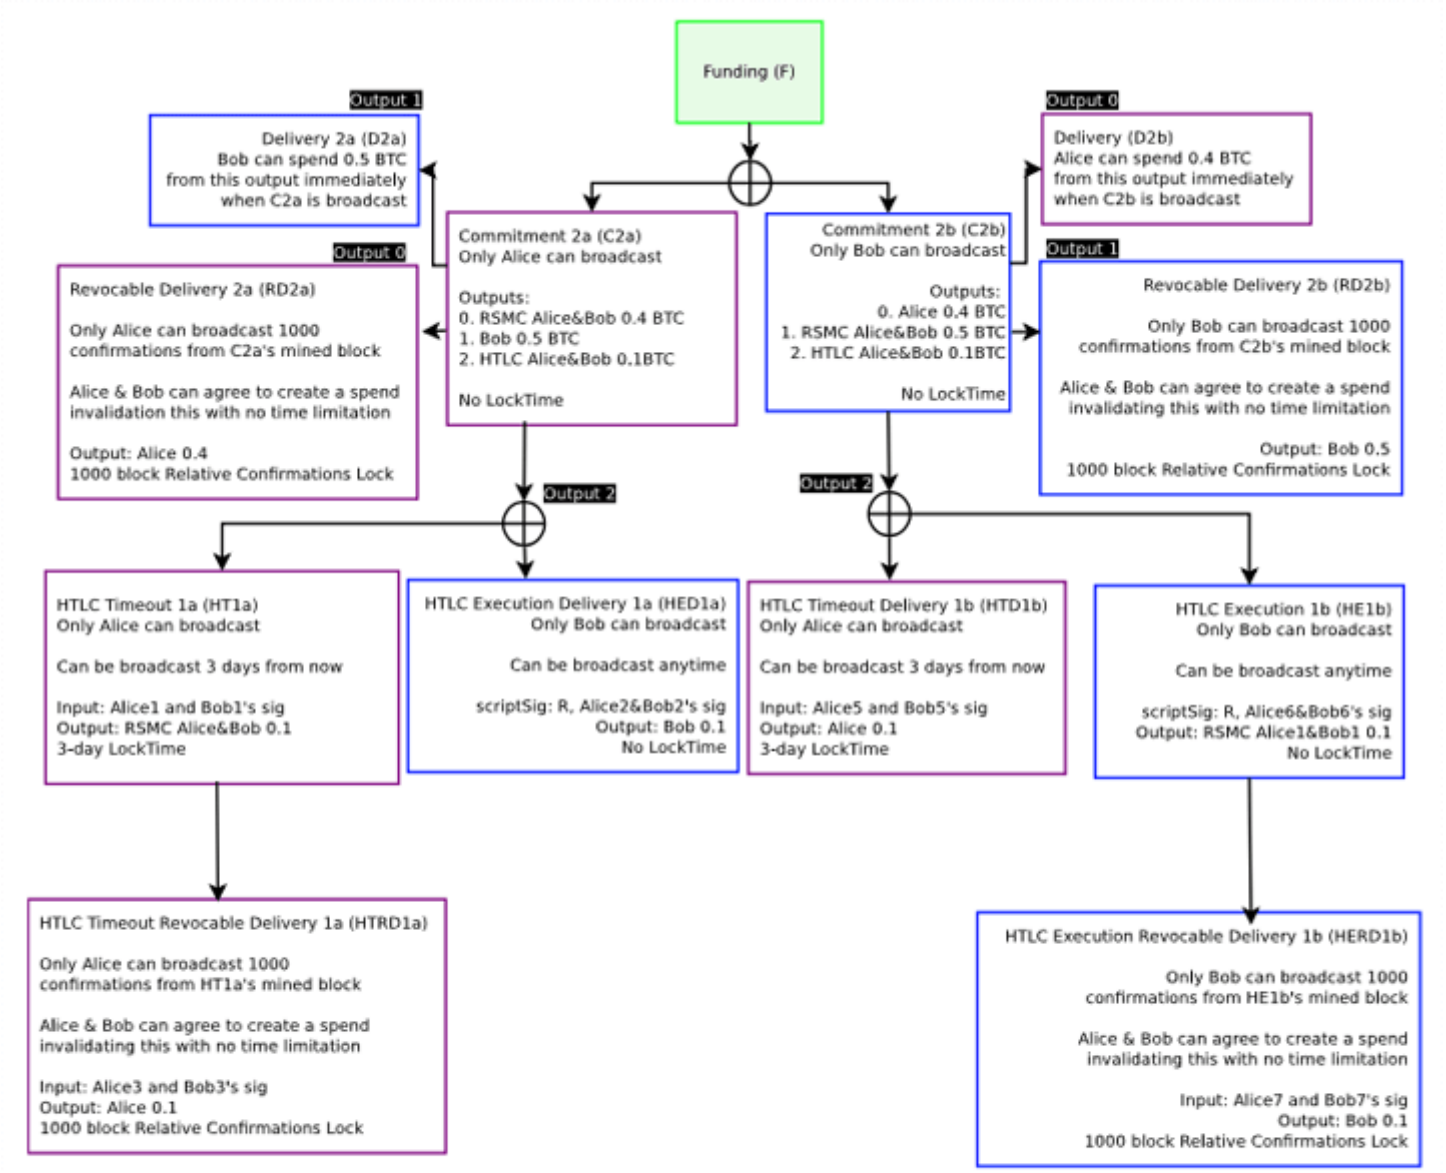

Figure 12: If Alice broadcasts C2a, then the left half will execute. If Bob broadcasts C2b, then the right half will execute. Either party may broadcast their Commitment transaction at any time. HTLC Timeout is only valid after 3 days. HTLC Executions can only be broadcast if the preimage to the hash R is known. Prior Commitments (and their dependent transactions) are not displayed for brevity.

Presume Alice and Bob wish to update their balance in the channel at Commitment 1 with a balance of 0.5 to

Alice and 0.5 to Bob.
Alice wishes to send 0.1 to Bob contingent upon knowledge of R within 3 days, after 3 days she wants her money back if Bob does not produce R.
The new Commitment Transaction will have a full refund of the current balance to Alice and Bob (Outputs 0 and 1), with output 2 being the HTLC, which describes the funds in transit. As 0.1 will be encumbered in an HTLC, Alice's balance is reduced to 0.4 and Bob's remains the same at 0.5.
This new Commitment Transaction (C2a/C2b) will have an HTLC output with two possible spends. Each spend is different depending on each counterparty's version of the Commitment Transaction. Similar to the bidirectional payment channel, when one party broadcasts their Commitment, payments to the counterparty will be assumed to be valid and not invalidated. This can occur because when one broadcasts a Commitment Transaction, one is attesting this is the most recent Commitment Transaction. If it is the most recent, then one is also attesting that the HTLC exists and was not invalidated before, so potential payments to one's counterparty should be valid.
Note that HTLC transaction names (beginning with the letter H) will begin with the number 1, whose values do not correlate with Commitment Transactions. This is simply the first HTLC transaction. HTLC transactions may persist between Commitment Transactions. Each HTLC has 4 keys per side of the transaction (C2a and C2b) for a total of 8 keys per counterparty.
The HTLC output in the Commitment Transaction has two sets of keys per counterparty in the output.
For Alice's Commitment Transaction (C2a), the HTLC output script requires multisig(PAlice2,PBob2) encumbered by disclosure of R, as well as multisig(PAlice1, PBob1) with no encumbering.
For Bob's Commitment Transaction (C2b), the HTLC output script requires multisig(PAlice6,PBob6) encumbered by disclosure of R, as well as multisig(PAlice5, PBob5) with no encumbering.
The HTLC output states are different depending upon which Commitment Transaction is broadcast.

HTLC when the Sender Broadcasts the Commitment Transaction

For the sender (Alice), the "Delivery" transaction is sent as an HTLC Execution Delivery transaction (HED1a), which is not encumbered in an RSMC. It assumes that this HTLC has never been terminated off-chain, as Alice is attesting that the broadcasted Commitment Transaction is the most recent. If Bob can produce the preimage R, he will be able to redeem funds from the HTLC after the Commitment Transaction is broadcast on the blockchain. This transaction consumes multisig(PAlice2,PBob2) if Alice broadcasts her Commitment C2a. Only Bob can broadcast HED1a since only Alice gave her signature for HED1a to Bob.
However, if 3 days have elapsed since forming the HTLC, then Alice will be able broadcast a "Timeout" transaction, the HTLC Timeout transaction (HT1a). This transaction is an RSMC. It consumes the output multisig(PAlice1,PBob1) without requiring disclosure of R if Alice broadcasts C2a. This transaction cannot enter into the blockchain until 3 days have elapsed. The output for this transaction is an RSMC with multisig(PAlice3,PBob3) with relative maturity of 1000 blocks, and multisig(PAlice4,PBob4) with no requirement for confirmation maturity. Only Alice can broadcast HT1a since only Bob gave his signature for HT1a to Alice.
After HT1a enters into the blockchain and 1000 block confirmations occur, an HTLC Timeout Revocable Delivery transaction (HTRD1a) may be broadcast by Alice which consumes multisig(PAlice3,PBob3). Only Alice can broadcast HTRD1a 1000 blocks after HT1a is broadcast since only Bob gave his signature for HTRD1a to Alice. This transaction can be revocable when another transaction supersedes HTRD1a using multisig(PAlice4,PBob4) which does not have any block maturity requirements.

HTLC when the Receiver Broadcasts the Commitment Transaction

For the potential receiver (Bob), the "Timeout" of receipt is refunded as an HTLC Timeout Delivery transaction (HTD1b). This transaction directly refunds the funds to the original sender (Alice) and is not encumbered in an RSMC. It assumes that this HTLC has never been terminated off-chain, as Bob is attesting that the broadcasted Commitment Transaction (C2b) is the most recent. If 3 days have elapsed, Alice can broadcast HTD1b and take the refund. This transaction consumes multisig(PAlice5,PAlice5) if Bob broadcasts C2b. Only Alice can broadcast HTD1b since Bob gave his signature for HTD1b to Alice.

However, if HTD1b is not broadcast (3 days have not elapsed) and Bob knows the preimage R, then Bob will be able to broadcast the HTLC Execution transaction (HE1b) if he can produce R. This transaction is an RSMC. It consumes the output multisig(PAlice6,PBob6) and requires disclosure of R if Bob broadcasts C2b. The output for this transaction is an RSMC with multisig(PAlice7,PBob7) with relative maturity of 1000 blocks, and multisig(PAlice8, PBob8) which does not have any block maturity requirements. Only Bob can broadcast HE1b since only Alice gave her signature for HE1b to Bob.

After HE1b enters into the blockchain and 1000 block confirmations occur, an HTLC Execution Revocable Delivery transaction (HERD1b) may be broadcast by Bob which consumes multisig(PAlice7,PBob7). Only Bob can broadcast HERD1b 1000 blocks after HE1b is broadcast since only Alice gave her signature for HERD1b to Bob. This transaction can be revocable when another transaction supersedes HERD1b using multisig(PAlice8,PBob8) which does not have any block maturity requirements.

HTLC Off-chain Termination

After an HTLC is constructed, to terminate an HTLC off-chain requires both parties to agree on the state of the channel. If the recipient can prove knowledge of R to the counterparty, the recipient is proving that they are able to immediately close out the channel on the Bitcoin blockchain and receive the funds. At this point, if both parties wish to keep the channel open, they should terminate the HTLC off-chain and create a new Commitment Transaction reflecting the new balance.

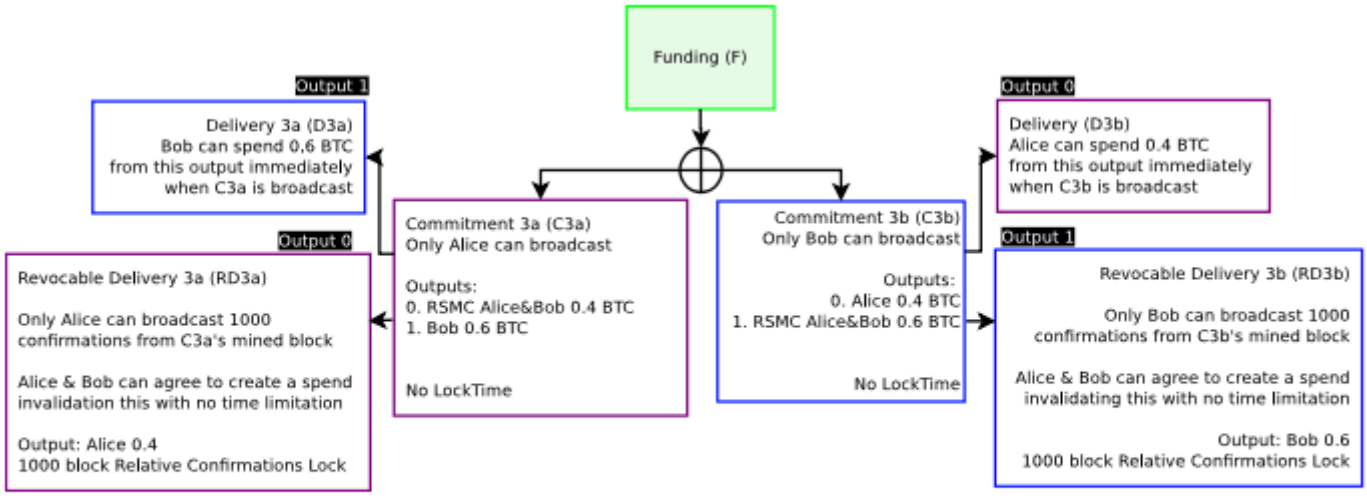

Figure 13: Since Bob proved to Alice he knows R by telling Alice R, Alice is willing to update the balance with a new Commitment Transaction. The payout will be the same whether C2 or C3 is broadcast at this time.

Similarly, if the recipient is not able to prove knowledge of R by disclosing R, both parties should agree to terminate the HTLC and create a new Commitment Transaction with the balance in the HTLC refunded to the sender.

If the counterparties cannot come to an agreement or become otherwise unresponsive, they should close out the channel by broadcasting the necessary channel transactions on the Bitcoin blockchain.

However, if they are cooperative, they can do so by first generating a new Commitment Transaction with the new balances, then invalidate the prior Commitment by exchanging Breach Remedy transactions (BR2a/BR2b). Additionally, if they are terminating a particular HTLC, they should also exchange some of their own private keys used in the HTLC transactions.

For example, Alice wishes to terminate the HTLC, Alice will disclose KAlice1 and KAlice4 to Bob. Correspondingly if Bob wishes to terminate the HTLC, Bob will disclose KBob6 and KBob8 to Alice. After the private keys are disclosed to the counterparty, if Alice broadcasts C2a, Bob will be able to take all the funds from the HTLC immediately. If Bob broadcasts C2b, Alice will be able to take all funds from the HTLC immediately. Note that when an HTLC is terminated, the older Commitment Transaction must be revoked as well.

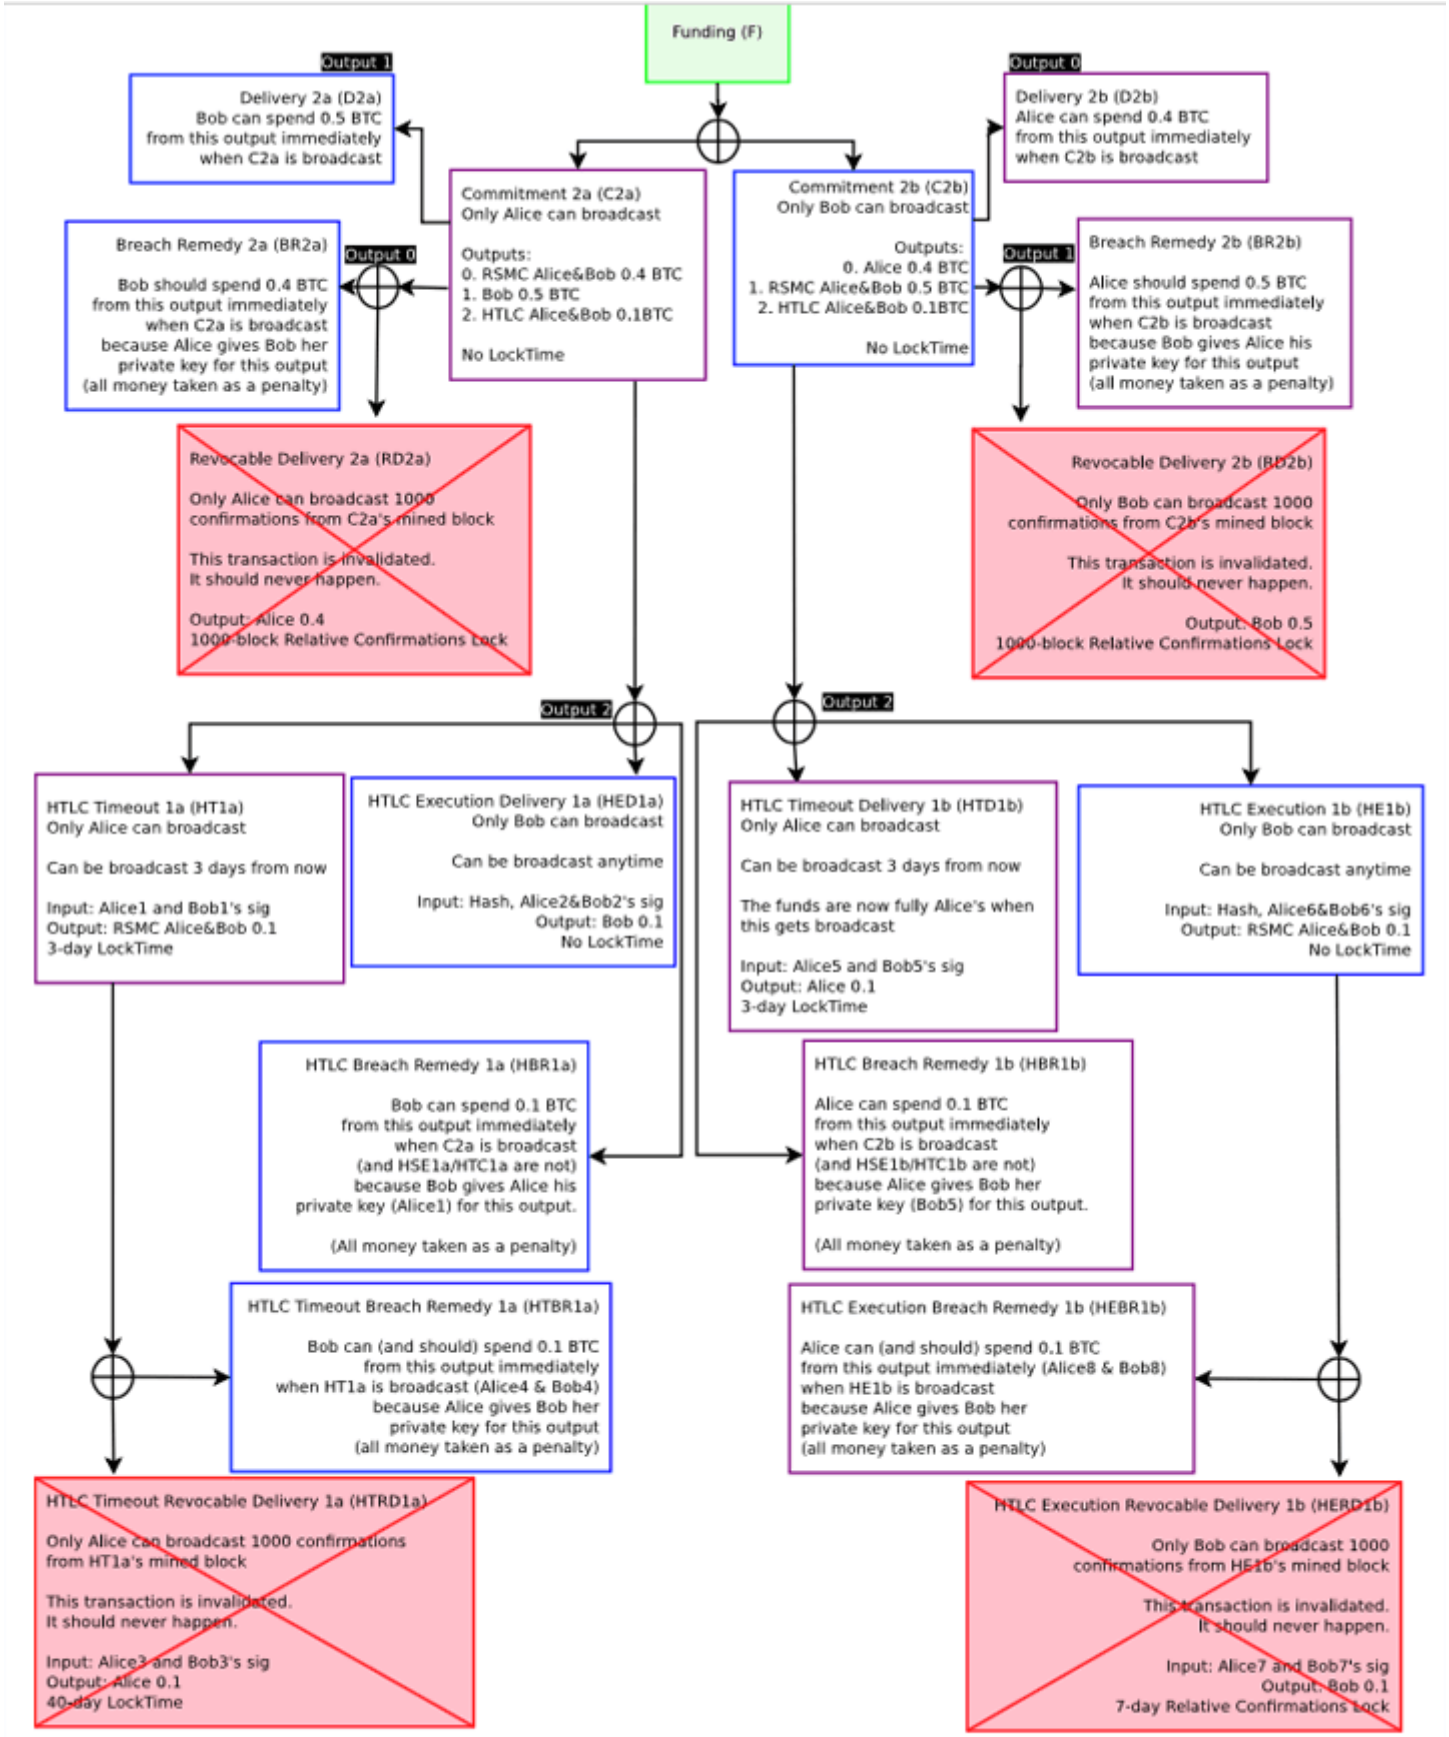

Figure 14: A fully revoked Commitment Transaction and terminated HTLC. If either party broadcasts Commitment 2, they will lose all their money to the counterparty. Other commitments (e.g. if Commitment 3 is the current Commitment) are not displayed for brevity.

Since both parties are able to prove the current state to each other, they can come to agreement on the current

balance inside the channel. Since they may broadcast the current state on the blockchain, they are able to come to agreement on netting out and terminating the HTLC with a new Commitment Transaction.

HTLC Formation and Closing Order

To create a new HTLC, it is the same process as creating a new Commitment Transaction, except the signatures for the HTLC are exchanged before the new Commitment Transaction's signatures.
To close out an HTLC, the process is as follows (from C2 to C3):

1. Alice signs and sends her signature for RD3b and C3b. At this point Bob can elect to broadcast C3b or C2b (with the HTLC) with the same payout. Bob is willing after receiving C3b to close out C2b.
2. Bob signs and sends his signature for RD3a and C3a, as well as his private keys used for Commitment 2 and the HTLC being terminated; he sends Alice KBobRSMC2, KBob5, and KBob8. At this point Bob should only broadcast C3b and should not broadcast C2b as he will lose all his money if he does so. Bob has fully revoked C2b and the HTLC. Alice is willing after receiving C3a to close out C2b.
3. Alice signs and sends her signature for RD3b and C3b, as well as her private keys used for Commitment 2 and the HTLC being terminated; she sends Bob KAliceRSMC2, KBob1, and KBob4. At this point neither party should broadcast Commitment 2, if they do so, their funds will be going to the counterparty. The old Commitment and old HTLC are now revoked and fully terminated. Only the new Commitment 3 remains, which does not have an HTLC.

When the HTLC has been closed, the funds are updated so that the present balance in the channel is what would occur had the HTLC contract been completed and broadcast on the blockchain. Instead, both parties elect to do off-chain novation and update their payments inside the channel.
It is absolutely necessary for both parties to complete off-chain novation within their designated time window. For the receiver (Bob), he must know R and update his balance with Alice within 3 days (or whatever time was selected), else Alice will be able to redeem it within 3 days. For Alice, very soon after her timeout becomes valid, she must novate or broadcast the HTLC Timeout transaction. She must also novate or broadcast the HTLC Timeout Revocable Delivery transaction as soon as it becomes valid. If the counterparty is unwilling to novate or is stalling, then one must broadcast the current channel state, including HTLC transactions) onto the Bitcoin blockchain.
The amount of time flexibility with these offers to novate are dependent upon one's contingent dependencies on the hashlock R. If one establishes a contract that the HTLC must be resolved within 1 day, then if the transaction times out Alice must resolve it by day 4 (3 days plus 1), else Alice risks losing funds.

Key Storage

Keys are generated using BIP 0032 Hierarchical Deterministic Wallets. Keys are pre-generated by both parties. Keys are generated in a merkle tree and are very deep within the tree. For instance, Alice pre-generates one million keys, each key being a child of the previous key. Alice allocates which keys to use according to some deterministic manner. For example, she starts with the child deepest in the tree to generate many sub-keys for day 1. This key is used as a master key for all keys generated on day 1. She gives Bob the address she wishes to use for the next transaction, and discloses the private key to Bob when it becomes invalidated. When Alice discloses to Bob all private keys derived from the day 1 master key and does not wish to continue using that master key, she can disclose the day 1 master key to Bob. At this point, Bob does not need to store all the keys derived from the day 1 master key. Bob does the same for Alice and gives her his day 1 key.
When all Day 2 private keys have been exchanged, for example by day 5, Alice discloses her Day 2 key. Bob

is able to generate the Day 1 key from the Day 2 key, as the Day 1 key is a child of the Day 2 key as well.
If a counterparty broadcasts the wrong Commitment Transaction, which private key to use in a transaction to recover funds can either be brute forced, or if both parties agree, they can use the sequence id number when creating the transaction to identify which sets of keys are used.
This enables participants in a channel to have prior output states (transactions) invalidated by both parties without using much data at all. By disclosing private keys pre-arranged in a merkle-tree, it is possible to invalidate millions of old transactions with only a few kilobytes of data per channel. Core channels in the zkFUND Network can conduct billions of transactions without a need for significant storage costs.

Blockchain Transaction Fees for Bidirectional Channels

It is possible for each participant to generate different versions of transactions to ascribe blame as to who broadcast the transaction on the blockchain. By having knowledge of who broadcast a transaction and the ability to ascribe blame, a third party service can be used to hold fees in a 2-of-3 multisig escrow. If one wishes to broadcast the transaction chain instead of agreeing to do a Funding Close or replacement with a new Commitment Transaction, one would communicate with the third party and broadcast the chain to the blockchain. If the counterparty refuses the notice from the third party to cooperate, the penalty is rewarded to the non-cooperative party. In most instances, participants may be indifferent to the transaction fees in the event of an uncooperative counterparty.
One should pick counterparties in the channel who will be cooperative, but is not an absolute necessity for the system to function. Note that this does not require trust among the rest of the network, and is only relevant for the comparatively minor transaction fees. The less trusted party may just be the one responsible for transaction fees.
The zkFUND Network fees will likely be significantly lower than blockchain transaction fees. The fees are largely derived from the time-value of locking up funds for a particular route, as well as paying for the chance of channel close on the blockchain. These should be significantly lower than on-chain transactions, as many transactions on a zkFUND Network channel can be settled into one single blockchain transaction. With a sufficiently robust and interconnected network, the fees should asymptotically approach negligibility for many types of transactions. With cheap fees and fast transactions, it will be possible to build scalable micropayments, even amongst high-frequency systems such as Internet of Things applications or per-unit micro-billing.

Pay to Contract

It is possible construct a cryptographically provable "Delivery Versus Payment" contract, or pay-to-contract, as proof of payment. This proof can be established as knowledge of the input R from hash(R) as payment of a certain value. By embedding a clause into the contract between the buyer and seller stating that knowing R is proof of funds sent, the recipient of funds has no incentive to disclose R unless they have certainty that they will receive payment. When the funds eventually get pulled from the buyer by their counterparty in their micropayment channel, R is disclosed as part of that pull of funds. One can design paper legal documents that specify that knowledge or disclosure of R implies fulfillment of payment. The sender can then arrange a cryptographically signed contract with knowledge of inputs for hashes treated as fulfillment of the paper contract before payment occurs.

The zkFUND Network

By having a micropayment channel with contracts encumbered by hashlocks and timelocks, it is possible to

clear transactions over a multi-hop payment network using a series of decrementing timelocks without additional central clearinghouses.

Traditionally, financial markets clear transactions by transferring the obligation for delivery at a central point and settle by transferring ownership through this central hub. Bank wire and fund transfer systems (such as ACH and the Visa card network), or equities clearinghouses (such as the DTCC) operate in this manner.

As Bitcoin enables programmatic money, it is possible to create transactions without contacting a central clearinghouse. Transactions can execute off-chain with no third party which collects all funds before disbursing it â€“ only transactions with uncooperative channel counterparties become automatically adjudicated on the blockchain.

The obligation to deliver funds to an end-recipient is achieved through a process of chained delegation. Each participant along the path assumes the obligation to deliver to a particular recipient. Each participant passes on this obligation to the next participant in the path. The obligation of each subsequent participant along the path, defined in their respective HTLCs, has a shorter time to completion compared to the prior participant. This way each participant is sure that they will be able to claim funds when the obligation is sent along the path.

Bitcoin Transaction Scripting, a form of what some call an implementation of "Smart Contracts", enables systems without trusted custodial clearinghouses or escrow services.

Decrementing Timelocks

Presume Alice wishes to send 0.001 BTC to Dave. She locates a route through Bob and Carol. The transfer path would be Alice to Bob to Carol to Dave.

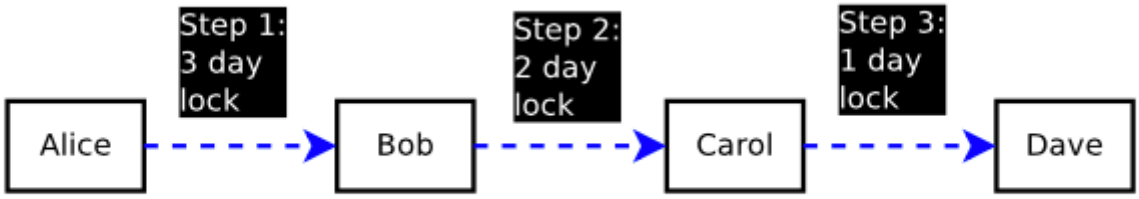

Figure 15: Payment over the Lightning Network using HTLCs.

When Alice sends payment to Dave through Bob and Carol, she requests from Dave hash(R) to use for this payment. Alice then counts the amount of hops until the recipient and uses that as the HTLC expiry. In this case, she sets the HTLC expiry at 3 days. Bob then creates an HTLC with Carol with an expiry of 2 days, and Carol does the same with Dave with an expiry of 1 day. Dave is now free to disclose R to Carol, and both parties will likely agree to immediate settlement via novation with a replacement Commitment Transaction. This then occurs step-by-step back to Alice. Note that this occurs off-chain, and nothing is broadcast to the blockchain when all parties are cooperative.

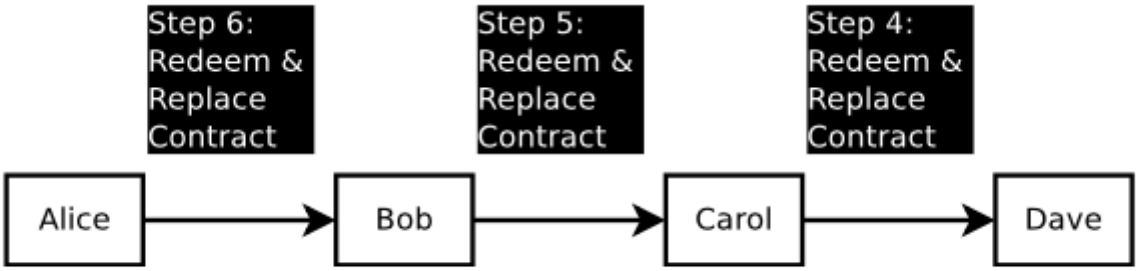

Figure 16: Settlement of HTLC, Alice's funds get sent to Dave.

Decrementing timelocks are used so that all parties along the path know that the disclosure of R will allow the disclosing party to pull funds, since they will at worst be pulling funds after the date whereby they must receive R. If Dave does not produce R within 1 day to Carol, then Carol will be able to close out the HTLC. If Dave broadcasts R after 1 day, then he will not be able to pull funds from Carol. Carol's responsibility to Bob occurs on day 2, so Carol will never be responsible for payment to Dave without an ability to pull funds from Bob provided that she updates her transaction with Dave via transmission to the blockchain or via novation.
In the event that R gets disclosed to the participants halfway through expiry along the path (e.g. day 2), then it is possible for some parties along the path to be enriched. The sender will be able to know R, so due to Pay to Contract, the payment will have been fulfilled even though the receiver did not receive the funds. Therefore, the receiver must never disclose R unless they have received an HTLC from their channel counterparty; they are guaranteed to receive payment from one of their channel counterparties upon disclosure of the preimage.
In the event a party outright disconnects, the counterparty will be responsible for broadcasting the current Commitment Transaction state in the channel to the blockchain. Only the failed non-responsive channel state gets closed out on the blockchain, all other channels should continue to update their Commitment Transactions via novation inside the channel. Therefore, counterparty risk for transaction fees are only exposed to direct channel counterparties. If a node along the path decides to become unresponsive, the participants not directly connected to that node suffer only decreased time-value of their funds by not conducting early settlement before the HTLC close.

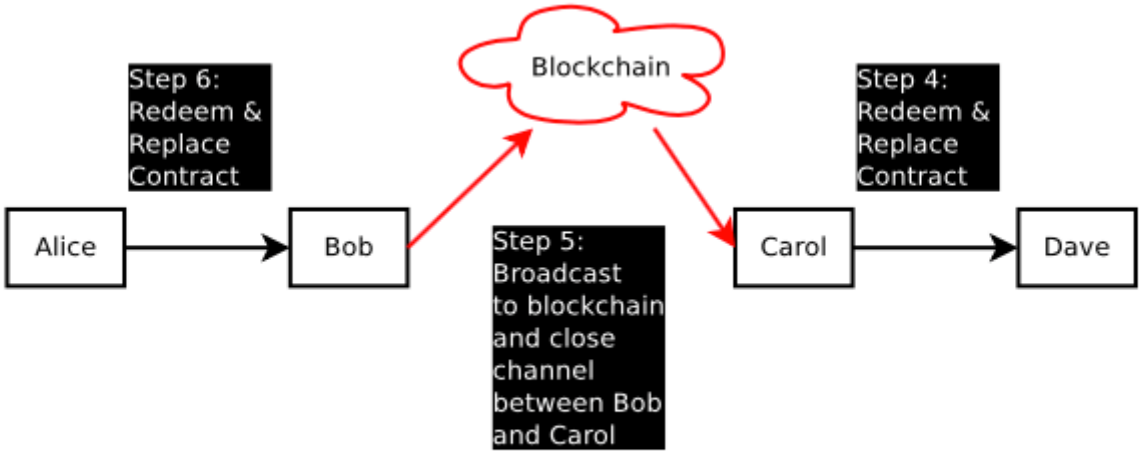

Figure 17: Only the non-responsive channels get broadcast on the blockchain, all others are settled off-chain via novation.

Payment Amount

It is preferable to use a small payment per HTLC. One should not use an extremely high payment, in case the payment does not fully route to its destination. If the payment does not reach its destination and one of the participants along the path is uncooperative, it is possible that the sender must wait until the expiry before receiving a refund. Delivery may be lossy, similar to packets on the internet, but the network cannot outright steal funds in transit. Since transactions don't hit the blockchain with cooperative channel counterparties, it is recommended to use as small of a payment as possible. A tradeoff exists between locking up transaction fees on each hop versus the desire to use as small a transaction amount as possible (the latter of which may incur higher total fees). Smaller transfers with more intermediaries imply a higher percentage paid as zkFUND Network fees to the intermediaries.

Clearing Failure and Rerouting

If a transaction fails to reach its final destination, the receiver should send an equal payment to the sender with the same hash, but not disclose R. This will net out the disclosure of the hash for the sender, but may not for the receiver. The receiver, who generated the hash, should discard R and never broadcast it. If one channel along the path cannot be contacted, then the channels may elect to wait until the path expires, which all participants will likely close out the HTLC as unsettled without any payment with a new Commitment Transaction.

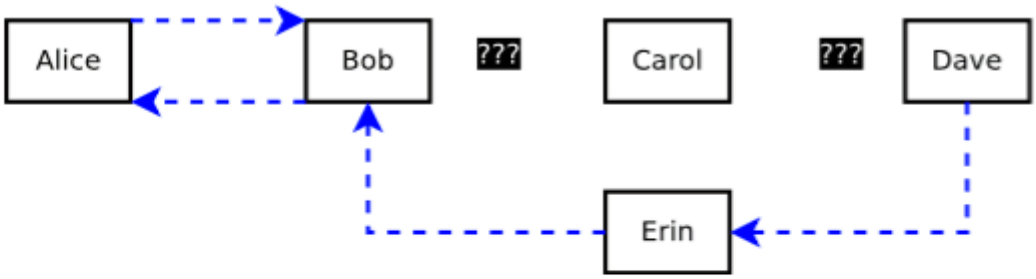

Figure 18: Dave creates a path back to Alice after Alice fails to send funds to Dave, because Carol is uncooperative. The input R from hash(R) is never brodcast by Dave, because Carol did not complete her actions. If R was broadcast, Alice will break-even. Dave, who controls R should never broadcast R because he may not receive funds from Carol, he should let the contracts expire. Alice and Bob have the option to net out and close the contract early, as well, in this diagram.

If the refund route is the same as the payment route, and there are no half-signed contracts whereby one party may be able to steal funds, it is possible to outright cancel the transaction by replacing it with a new Commitment Transaction starting with the most recent node who participated in the HTLC.

It is also possible to clear out a channel by creating an alternate route path in which payment will occur in the opposite direction (netting out to zero) and/or creating an entirely alternate route for the payment path. This will create a time-value of money for disclosing inputs to hashes on the zkFUND Network. Participants may specialize in high connectivity between nodes and offering to offload contract hashlocks from other nodes for a fee. These participants will agree to payments which net out to zero (plus fees), but are loaning bitcoins for a set time period. Most likely, these entities with low demand for channel resources will be end-users who are already connected to multiple well-connected nodes. When an end-user connects to a node, the node may ask the client to lock up their funds for several days to another channel the client has established for a fee. This can be achieved by having the new transactions require a new hash(Y) from input Y in addition to the existing hash which may be generated by any participant, but must disclose Y only after a full circle is established. The new participant has the same responsibility as well as the same timelocks as the old participant being replaced. It is also possible that the one new participant replaces multiple hops.

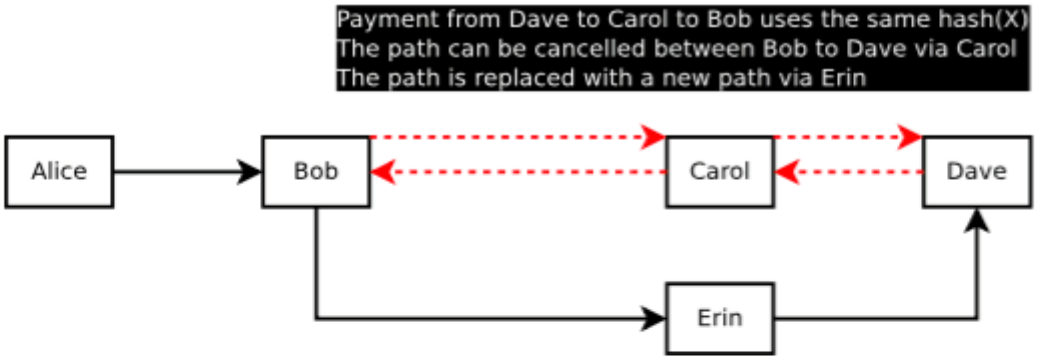

Figure 19: Erin is connected to both Bob and Dave. If Bob wishes to free up his channel with Carol, since that channel is active and very profitable, Bob can offload the payment to Dave via Erin. Since Erin has extra bitcoin available, she will be able to collect some fee for offloading the channel between Bob and Carol as well as between Carol and Dave. The channels between Bob and Carol as well as Carol and Dave are undone and no longer have the HTLC, nor has payment occurred on that path. Payment will occur on the path involving Erin. This is achieved by creating a new payment from Dave to Carol to Bob contingent upon Erin constructing an HTLC. The payment in dashed lines (red) are netted out to zero and settled via a new Commitment Contract.

Payment Routing

It is theoretically possible to build a route map implicitly from observing 2-of-2 multisigs on the blockchain to build a routing table. Note, however, this is not feasible with pay-to-script-hash transaction outputs, which can be resolved out-of-band from the bitcoin protocol via a third party routing service. Building a routing table will become necessary for large operators (e.g. BGP, Cjdns). Eventually, with optimizations, the network will look a lot like the correspondent banking network, or Tier-1 ISPs. Similar to how packets still reach their destination on your home network connection, not all participants need to have a full routing table. The core Tier-1 routes can be online all the time -while nodes at the edges, such as average users, would be connected intermittently.

Node discovery can occur along the edges by pre-selecting and offering partial routes to well-known nodes.

Correctness

In order to achieve correctness, given a maximal amount of Byzantine failures, it must be shown that it is impossible for a fraudulent transaction to be confirmed during consensus, unless the number of faulty nodes exceeds that tolerance. The proof of the correctness of the RPCA then follows directly: since a transaction is only approved if 80% of the UNL of a server agrees with it, as long as 80% of the UNL is honest, no fraudulent transactions will be approved. Thus for a UNL of n nodes in the network, the consensus protocol will maintain correctness so long as:

$$f \leq (n-1)/5$$

where f is the number Byzantine failures. In fact, even in the face of (n - 1)/5 + 1 Byzantine failures, correctness is still technically maintained. The consensus process will fail, but it will still not be possible to confirm a fraudulent transaction. Indeed it would take (4n + 1)/5 Byzantine failures for an incorrect transaction to be confirmed. We call this second bound the bound for weak correctness, and the former the bound for strong correctness.

It should also be noted that not all "fraudulent" trans- actions pose a threat, even if confirmed during consensus. Should a user attempt to double-spend funds in two transactions, for example, even if both transactions are confirmed during the consensus process, after the first transaction is applied, the second will fail, as the funds are no longer available. This robustness is due to the fact that transactions are applied deterministically, and that consensus ensures that all nodes in the network are applying the deterministic rules to the same set of transactions.

For a slightly different analysis, let us assume that the probability that any node will decide to collude and join a nefarious cartel is pc. Then the probability of correctness is given by p*, where:

$$p^* = \sum_{i=0}^{\lceil (\frac{n-1}{5}) \rceil} \binom{n}{i} p_c^i (1-p_c)^{n-i}$$

This probability represents the likelihood that the size of the nefarious cartel will remain below the maximal threshold of Byzantine failures, given pc. Since this likelihood is a binomial distribution, values of pc greater than 20% will result in expected cartels of size greater than 20% of the network, thwarting the consensus process. In practice, a UNL is not chosen randomly, but rather with the intent to minimize pc. Since nodes are not anonymous but rather cryptographically identifiable, selecting a UNL of nodes from a mixture of continents, nations, industries, ideologies, etc. will produce values of pc much lower than 20%. As an example, the probability of the Anti-Defamation League and the Westboro Baptist Church colluding to defraud the network, is certainly much, much smaller than 20%. Even if the UNL has a relatively large pc, say 15%, the probability of correctness is extremely high even with only 200 nodes in the UNL: 97.8%.

A graphical representation of how the probability of incorrectness scales as a function of UNL size for differing values of pc is depicted in Figure 1. Note that here the vertical axis represents the probability of a nefarious cartel thwarting consensus, and thus lower values indi- cate greater probability of consensus success. As can be seen in the figure, even with a pc as high as 10%, the probability of consensus being thwarted very quickly becomes negligible as the UNL grows past 100 nodes.

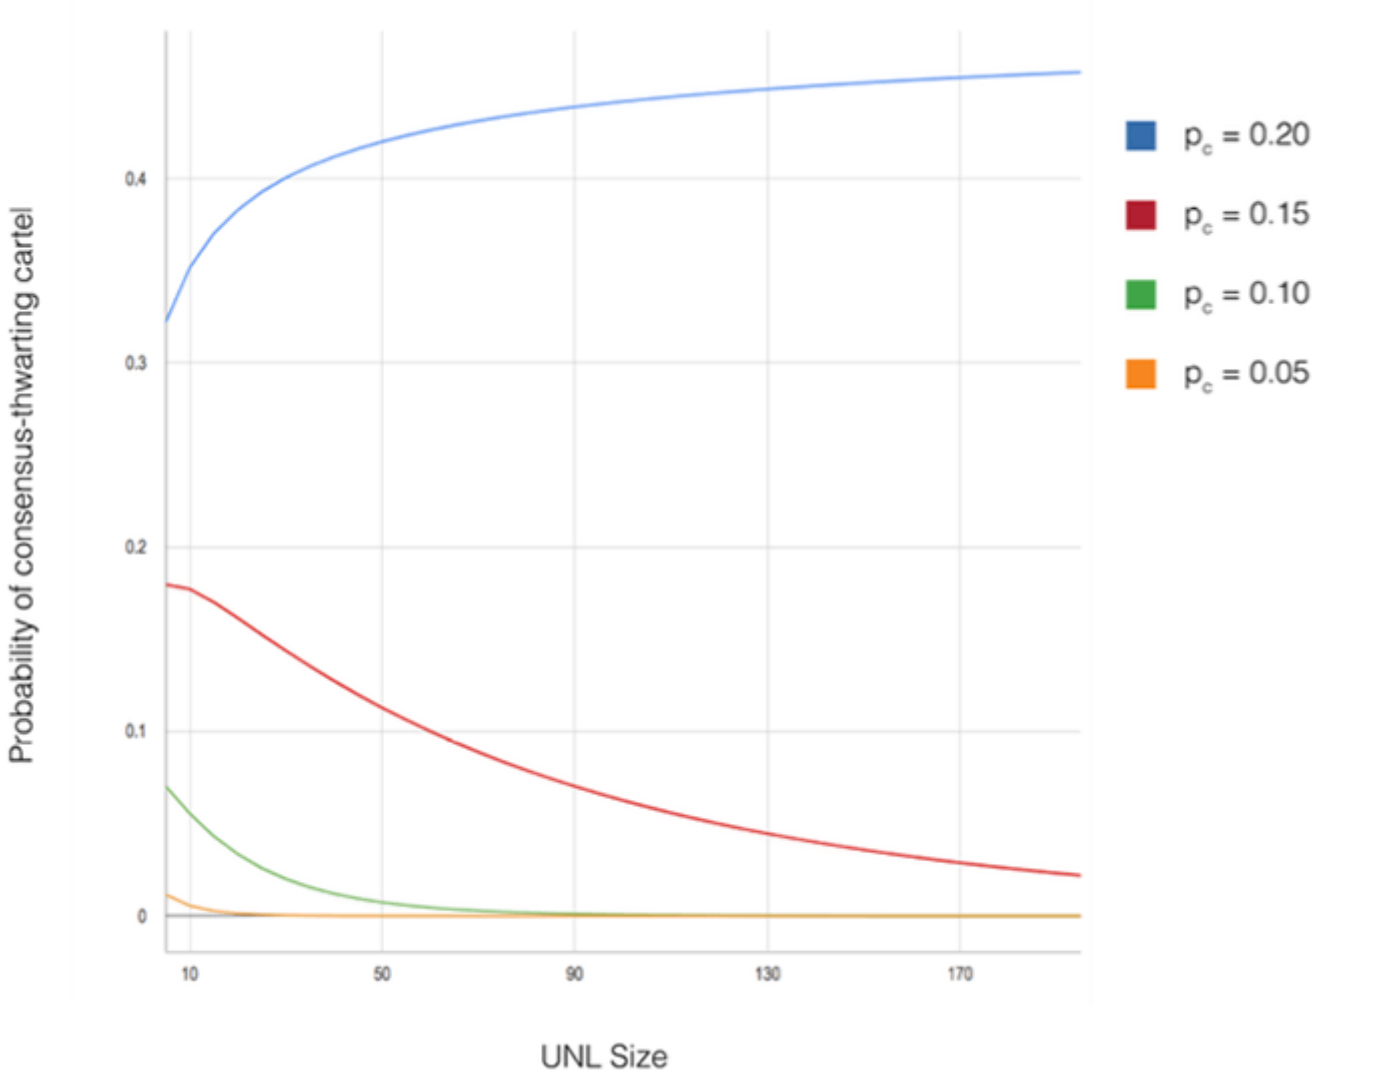

Figure 1. Probability of a nefarious cartel being able to thwart consensus as a function of the size of the UNL, for different values of pc, the probability that any member of the UNL will decide to collude with others. Here, lower values indicate a higher probability of consensus success.

Agreement

To satisfy the agreement requirement, it must be shown that all nonfaulty nodes reach consensus on the same set of transactions, regardless of their UNLs. Since the UNLs for each server can be different, agreement is not inherently guaranteed by the correctness proof. For example, if there are no restrictions on the membership of the UNL, and the size of the UNL is not larger than 0.2 * ntotal where ntotal is the number of nodes in the entire network, then a fork is possible. This is illustrated by a simple example (depicted in figure 2): imagine two cliques within the UNL graph, each larger than 0.2 * ntotal. By cliques, we mean a set of nodes where each nodeâ€™s UNL is the selfsame set of nodes. Because these two cliques do not share any members, it is possible for each to achieve a correct consensus independently of each other, violating agreement. If the connectivity of the two cliques surpasses 0.2 * ntotal, then a fork is no longer possible, as disagreement between the cliques would prevent consensus from being reached at the 80% agreement threshold that is required.

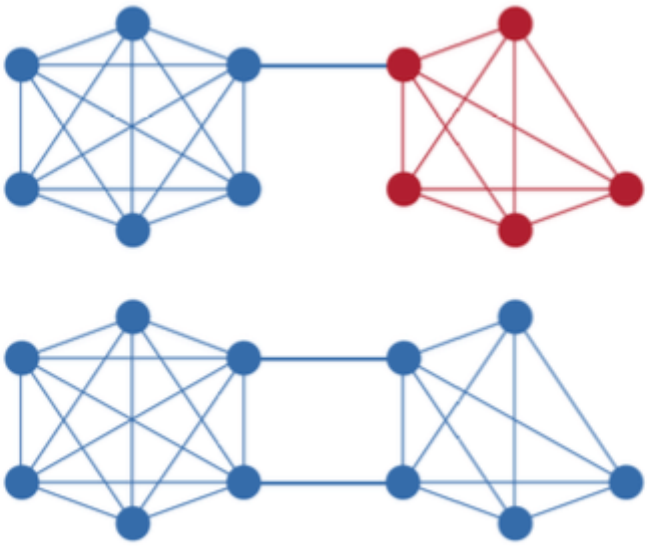

Figure 2. An example of the connectivity required to prevent a fork between two UNL cliques.

An upper bound on the connectivity required to prove agreement is given by:

$$|UNL_i \cap UNL_j| \geq \frac{1}{5}\max(|UNL_i|, |UNL_j|) \forall i, j$$

This upper bound assumes a clique-like structure of UNLs, i.e. nodes form sets whose UNLs contain other nodes in those sets. This upper bound guarantees that no two cliques can reach consensus on conflicting transactions, since it becomes impossible to reach the 80% threshold required for consensus. A tighter bound is possible when indirect edges between UNLs are taken into account as well. For example, if the structure of the network is not clique-like, a fork becomes much more difficult to achieve, due to the greater entanglement of the UNLs of all nodes.

It is interesting to note that no assumptions are made about the nature of the intersecting nodes. The intersection of two UNLs may include faulty nodes, but so long as the size of the intersection is larger than the bound required to guarantee agreement, and the total number of faulty nodes is less than the bound required to satisfy strong correctness, then both correctness and agreement will be achieved. That is to say, agreement is dependent solely on the size of the intersection of nodes, not on the size of the intersection of nonfaulty nodes.

Utility

While many components of utility are subjective, one that is indeed provable is convergence: that the consensus process will terminate in finite time.

We define convergence as the point in which the RPCA reaches consensus with strong correctness on the ledger, and that ledger then becomes the last-closed ledger. Note that while technically weak correctness still represents convergence of the algorithm, it is only convergence in the trivial case, as proposition C3 is violated, and no transactions will ever be confirmed. From the results above, we know that strong correctness is always achievable in the face of up to (n âˆ’ 1)/5 Byzantine failures, and that only one consensus will be achieved in the entire network so long as the UNL-connectedness con- dition is met (Equation 3). All that

remains is to show that when both of these conditions are met, consensus is reached in finite time.
Since the consensus algorithm itself is deterministic, and has a preset number of rounds, t, before consensus is terminated, and the current set of transactions are declared approved or not-approved (even if at this point no transactions have more than the 80% required agreement, and the consensus is only the trivial consensus), the limiting factor for the termination of the algorithm is the communication latency between nodes. In order to bound this quantity, the response-time of nodes is monitored, and nodes whoâ€™s latency grows larger than a preset bound b are removed from all UNLs. While this guarantees that consensus will terminate with an upper bound of tb, it is important to note that the bounds described for correctness and agreement above must be met by the final UNL, after all nodes that will be dropped have been dropped. If the conditions hold for the initial UNLs for all nodes, but then some nodes are dropped from the network due to latency, the correctness and agreement guarantees do not automatically hold but must be satisfied by the new set of UNLs.
As mentioned above, a latency bound heuristic is en- forced on all nodes in the zkFUND Network to guarantee that the consensus algorithm will converge. In addi- tion, there are a few other heuristics and procedures that provide utility to the RPCA.

There is a mandatory 2 second window for all nodes to propose their initial candidate sets in each round of consensus. While this does introduce a lower bound of 2 seconds to each consensus round, it also guarantees that all nodes with reasonable latency will have the ability to participate in the consensus process.

As the votes are recorded in the ledger for each round of consensus, nodes can be flagged and removed from the network for some common, easily-identifiable malicious behaviors. These include nodes that vote "No" on every transaction, and nodes that consistently propose transactions which are not validated by consensus.

A curated default UNL is provided to all users, which is chosen to minimize pc, described in section 3.2. While users can and should select their own UNLs, this default list of nodes guarantees that even naive users will participate in a consensus process that achieves correctness and agreement with extremely high probability.

A network split detection algorithm is also employed to avoid a fork in the network. While the consensus algorithm certifies that the transactions on the last-closed ledger are correct, it does not prohibit the possibility of more than one last-closed ledger existing on different subsections of the network with poor connectivity. To try and identify if such a split has occurred, each node monitors the size of the active members of its UNL. If this size suddenly drops below a preset threshold, it is possible that a split has occurred. In order to prevent a false

```
positive in the case where a large section of a UNL has temporary latency,
nodes are allowed to publish a "partial validation", in which they do not
process or vote on transactions, but declare that are still participating
in the consensus process, as opposed to a different consensus process on a
disconnected subnetwork.
```

```
While it would be possible to apply the RPCA in just one round of
consensus, utility can be gained through multiple rounds, each with an
increasing minimum-required percentage of agreement, before the final
round with an 80% requirement. These rounds allow for detection of latent
nodes in the case that a few such nodes are creating a bottleneck in the
transaction rate of the network. These nodes will be able to initially
keep up during the lower-requirement rounds but fall behind and be
identified as the threshold increases. In the case of one round of
consensus, it may be the case that so few transactions pass the 80%
threshold, that even slow nodes can keep up, lowering the transaction rate
of the entire network.
```

Simulation Code

The provided simulation code demonstrates a round of RPCA, with parameterizable features (the number of nodes in the network, the number of malicious nodes, latency of messages, etc.). The simulator begins in perfect disagreement (half of the nodes in the network initially propose "yes", while the other half propose "no"), then proceeds with the consensus process, showing at each stage the number of yes/no votes in the network as nodes adjust their proposals based upon the proposals of their UNL members. Once the 80% threshold is reached, consensus is achieved. We encourage the reader to experiment with different values of the constants defined at the beginning of "Sim.cpp", in order to become familiar with the consensus process under different conditions.

Risks

The primary risks relate to timelock expiration. Additionally, for core nodes and possibly some merchants to be able to route funds, the keys must be held online for lower latency. However, end-users and nodes are able to keep their private keys firewalled off in cold storage.

Improper Timelocks

Participants must choose timelocks with sufficient amounts of time. If insufficient time is given, it is possible that timelocked transactions believed to be invalid will become valid, enabling coin theft by the counterparty. There is a trade-off between longer timelocks and the time-value of money. When writing wallet and zkFUND Network application software, it is necessary to ensure that sufficient time is given and users are able to have their transactions enter into the blockchain when interacting with non-cooperative or malicious channel counterparties.

Forced Expiration Spam

Forced expiration of many transactions may be the greatest systemic risk when using the zkFUND Network. If a malicious participant creates many channels and forces them all to expire at once, these may overwhelm block data capacity, forcing expiration and broadcast to the blockchain. The result would be mass spam on the bitcoin network. The spam may delay transactions to the point where other locktimed transactions become valid.

This may be mitigated by permitting one transaction replacement on all pending transactions. Anti-spam can be used by permitting only one transaction replacement of a higher sequence number by the inverse of an even or odd number. For example, if an odd sequence number was broadcast, permit a replacement to a higher even number only once. Transactions would use the sequence number in an orderly way to replace other transactions. This mitigates the risk assuming honest miners. This attack is extremely high risk, as incorrect broadcast of Commitment Transactions entail a full penalty of all funds in the channel.

Additionally, one may attempt to steal HTLC transactions by forcing a timeout transaction to go through when it should not. This can be easily mitigated by having each transfer inside the channel be lower than the total transaction fees used. Since transactions are extremely cheap and do not hit the blockchain with cooperative channel counterparties, large transfers of value can be split into many small transfers. This attempt can only work if the blocks are completely full for a long time. While it is possible to mitigate it using a longer HTLC timeout duration, variable block sizes may become common, which may need mitigations.

If this type of transaction becomes the dominant form of transactions which are included on the blockchain, it may become necessary to increase the block size and run a variable blocksize structure and timestop flags as described in the section below. This can create sufficient penalties and disincentives to be highly unprofitable and unsuccessful for attackers, as attackers lose all their funds from broadcasting the wrong transaction, to the point where it will never occur.

Coin Theft via Cracking

As parties must be online and using private keys to sign, there is a possibility that, if the computer where the private keys are stored is compromised, coins will be stolen by the attacker. While there may be methods to mitigate the threat for the sender and the receiver, the intermediary nodes must be online and will likely be processing the transaction automatically. For this reason, the intermediary nodes will be at risk and should not be holding a substantial amount of money in this "hot wallet." Intermediary nodes which have better security will likely be able to out-compete others in the long run and be able to conduct greater transaction volume due to lower fees. Historically, one of the largest component of fees and interest in the financial system are from various forms of counterparty risk â€“ in Bitcoin it is possible that the largest component in fees will be derived from security risk premiums.

A Funding Transaction may have multiple outputs with multiple Commitment Transactions, with the Funding Transaction key and some Commitment Transactions keys stored offline. It is possible to create an equivalent of a "Checking Account" and "Savings Account" by moving funds between outputs from a Funding Transaction, with the "Savings Account" stored offline and requiring additional signatures from security services.

Data Loss

When one party loses data, it is possible for the counterparty to steal funds. This can be mitigated by having a third party data storage service where encrypted data gets sent to this third party service which the party cannot decrypt. Additionally, one should choose channel counterparties who are responsible and willing to provide the current state, with some periodic tests of honesty.

Forgetting to Broadcast the Transaction in Time

If one does not broadcast a transaction at the correct time, the counterparty may steal funds. This can be mitigated by having a designated third party to send funds. An output fee can be added to create an incentive for this third party to watch the network. Further, this can also be mitigated by implementing OP CHECKSEQUENCEVERIFY.

Inability to Make Necessary Soft-Forks

Changes are necessary to bitcoin, such as the malleability soft-fork. Additionally, if this system becomes popular, it will be necessary for the system to securely transact with many users and some kind of structure like a blockheight timestop will be desirable. This system assumes such changes to enable zkFUND Network to exist entirely, as well as soft-forks ensuring the security is robust against attackers will occur. While the system may continue to operate with only some time lock and malleability soft-forks, there will be necessary soft-forks regarding systemic risks. Without proper community foresight, an inability to establish a timestop or similar function will allow systemic attacks to take place and may not be recognized as imperative until an attack actually occurs.

Colluding Miner Attacks

Miners may elect to refuse to enter in particular transactions (e.g. Breach Remedy transactions) in order to assist in timeout coin theft. An attacker can pay off all miners to refuse to include certain transactions in their mempool and blocks. The miners can identify their own blocks in an attempt to prove their behavior to the paying attacker.

This can be mitigated by encouraging miners to avoid identifying their own blocks. Further, it should be expected that this kind of payment to miners is malicious activity and the contract is unenforcible. Miners may then take payment and surreptitiously mine a block without identifying the block to the attacker. Since the attacker is paying for this, they will quickly run out of money by losing the fee to the miner, as well as losing all their money in the channel. This attack is unlikely and fairly unattractive as it is far too difficult and requires a high degree of collusion with extreme risk.

The risk model of this attack occurirng is similar to that of miners colluding to do reorg attacks: Extremely unlikely with many uncoordinated miners.

Block Size Increases and Consensus

If we presume that a decentralized payment network exists and one user will make 3 blockchain transactions per year on average, Bitcoin will be able to support over 35 million users with 1MB blocks in ideal circumstances (assuming 2000 transactions/MB, or 500 bytes/Tx). This is quite limited, and an increase of the block size may be necessary to support everyone in the world using Bitcoin. A simple increase of the block size would be a hard fork, meaning all nodes will need to update their wallets if they wish to participate in the network with the larger blocks.

While it may appear as though this system will mitigate the block size increases in the short term, if it achieves global scale, it will necessitate a block size increase in the long term. Creating a credible tool to help prevent blockchain spam designed to encourage transactions to timeout becomes imperative.

To mitigate timelock spam vulnerabilities, non-miner and miners' consensus rules may also differ if the miners' consensus rules are more restrictive. Non-miners may accept blocks over 1MB, while miners may have different soft-caps on block sizes. If a block size is above that cap, then that is viewed as an invalid block by other miners, but not by non-miners. The miners will only build the chain on blocks which are valid

according to the agreed-upon soft-cap. This permits miners to agree on raising the block size limit without requiring frequent hard-forks from clients, so long as the amount raised by miners does not go over the clients' hard limit. This mitigates the risk of mass expiry of transactions at once. All transactions which are not redeemed via Exercise Settlement (ES) may have a very high fee attached, and miners may use a consensus rule whereby those transactions are exempted from the soft-cap, making it very likely the correct transactions will enter the blockchain.
When transactions are viewed as circuits and contracts instead of transaction packets, the consensus risks can be measured by the amount of time available to cover the UTXO set controlled by hostile parties. In effect, the upper bound of the UTXO size is determined by transaction fees and the standard minimum transaction output value. If the bitcoin miners have a deterministic mempool which prioritizes transactions respecting a "weak" local time order of transactions, it could become extremely unprofitable and unlikely for an attack to succeed. Any transaction spam time attack by broadcasting the incorrect Commitment Transaction is extremely high risk for the attacker, as it requires an immense amount of bitcoin and all funds committed in those transactions will be lost if the attacker fails.

Use Cases

In addition to helping bitcoin scale, there are many uses for transactions on the zkFUND Network:

```
Instant Transactions. Using zkFUND, Bitcoin transactions are now nearly
instant with any party. It is possible to pay for a cup of coffee with
direct non-revocable payment in milliseconds to seconds.
```

```
Exchange Arbitrage. There is presently incentive to hold funds on
exchanges to be ready for large market moves due to 3-6 block confirmation
times. It is possible for the exchange to participate in this network and
for clients to move their funds on and off the exchange for orders nearly
instantly. If the exchange does not have deep market depth and commits to
only permitting limit orders close to the top of the order book, then the
risk of coin theft becomes much lower. The exchange, in effect, would no
longer have any need for a cold storage wallet. This may substantially
reduce thefts and the need for trusted third party custodians.
```

```
Micropayments. Bitcoin blockchain fees are far too high to accept
micropayments, especially with the smallest of values. With this system,
near-instant micropayments using Bitcoin without a 3rd party custodian
would be possible. It would enable, for example, paying per-megabyte for
internet service or per-article to read a newspaper.
```

```
Financial Smart Contracts and Escrow. Financial contracts are especially
time-sensitive and have higher demands on blockchain computation. By
moving the overwhelming majority of trustless transactions off-chain, it
```

is possible to have highly complex transaction contract terms without ever hitting the blockchain.

Cross-Chain Payments. So long as there are similar hash-functions across chains, it's possible for transactions to be routed over multiple chains with different consensus rules. The sender does not have to trust or even know about the other chains â€“ even the destination chain. Simiarly, the receiver does not have to know anything about the sender's chain or any other chain. All the receiver cares about is a conditional payment upon knowledge of a secret on their chain. Payment can be routed by participants in both chains in the hop. E.g. Alice is on Bitcoin, Bob is on both Bitcoin and X-Coin and Carol is on a hypothetical X-Coin, Alice can pay Carol without understanding the X-Coin consensus rules.

Calculations

We consider the scenario of an attacker trying to generate an alternate chain faster than the honest chain. Even if this is accomplished, it does not throw the system open to arbitrary changes, such as creating value out of thin air or taking money that never belonged to the attacker. Nodes are not going to accept an invalid transaction as payment, and honest nodes will never accept a block containing them. An attacker can only try to change one of his own transactions to take back money he recently spent.

The race between the honest chain and an attacker chain can be characterized as a Binomial Random Walk. The success event is the honest chain being extended by one block, increasing its lead by +1, and the failure event is the attacker's chain being extended by one block, reducing the gap by -1.

The probability of an attacker catching up from a given deficit is analogous to a Gambler's Ruin problem. Suppose a gambler with unlimited credit starts at a deficit and plays potentially an infinite number of trials to try to reach breakeven. We can calculate the probability he ever reaches breakeven, or that an attacker ever catches up with the honest chain, as follows:

p = probability an honest node finds the next block
q = probability the attacker finds the next block
q_z = probability the attacker will ever catch up from z blocks behind

$$q_z = \begin{cases} 1 & \text{if } p \leq q \\ (q/p)^z & \text{if } p > q \end{cases}$$

Given our assumption that p > q, the probability drops exponentially as the number of blocks the attacker has to catch up with increases. With the odds against him, if he doesn't make a lucky lunge forward early on, his chances become vanishingly small as he falls further behind.

We now consider how long the recipient of a new transaction needs to wait before being sufficiently certain the sender can't change the transaction. We assume the sender is an attacker who wants to make the recipient believe he paid him for a while, then switch it to pay back to himself after some time has passed. The receiver will be alerted when that happens, but the sender hopes it will be too late.

The receiver generates a new key pair and gives the public key to the sender shortly before signing. This prevents the sender from preparing a chain of blocks ahead of time by working on it continuously until he is lucky enough to get far enough ahead, then executing the transaction at that moment. Once the transaction is sent, the dishonest sender starts working in secret on a parallel chain containing an alternate version of his transaction.

The recipient waits until the transaction has been added to a block and z blocks have been linked after it. He doesn't know the exact amount of progress the attacker has made, but assuming the honest blocks took the average expected time per block, the attacker's potential progress will be a Poisson distribution with expected value:

$$\lambda = z\frac{q}{p}$$

To get the probability the attacker could still catch up now, we multiply the Poisson density for each amount of progress he could have made by the probability he could catch up from that point:

$$\sum_{k=0}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} \cdot \begin{cases} (q/p)^{(z-k)} & \text{if } k \le z \\ 1 & \text{if } k > z \end{cases}$$

Rearranging to avoid summing the infinite tail of the distribution...

$$1 - \sum_{k=0}^{z} \frac{\lambda^k e^{-\lambda}}{k!} \left(1 - (q/p)^{(z-k)}\right)$$

Converting to C code...

```
#include <math.h>
double AttackerSuccessProbability(double q, int z)
{
    double p = 1.0 - q;
    double lambda = z * (q / p);
    double sum = 1.0;
    int i, k;
    for (k = 0; k <= z; k++)
    {
        double poisson = exp(-lambda);
        for (i = 1; i <= k; i++)
            poisson *= lambda / i;
        sum -= poisson * (1 - pow(q / p, z - k));
    }
    return sum;
}
```

Running some results, we can see the probability drop off exponentially with z.

```
q=0.1
z=0     P=1.0000000
z=1     P=0.2045873
z=2     P=0.0509779
z=3     P=0.0131722
z=4     P=0.0034552
z=5     P=0.0009137
z=6     P=0.0002428
z=7     P=0.0000647
z=8     P=0.0000173
z=9     P=0.0000046
z=10    P=0.0000012

q=0.3
z=0     P=1.0000000
z=5     P=0.1773523
z=10    P=0.0416605
z=15    P=0.0101008
z=20    P=0.0024804
z=25    P=0.0006132
z=30    P=0.0001522
z=35    P=0.0000379
z=40    P=0.0000095
z=45    P=0.0000024
z=50    P=0.0000006
```

Solving for P less than 0.1%...

```
P < 0.001
q=0.10   z=5
q=0.15   z=8
q=0.20   z=11
q=0.25   z=15
q=0.30   z=24
q=0.35   z=41
q=0.40   z=89
q=0.45   z=340
```

Related works

The original Bitcoin proof-of-work protocol uses the CPU-intensive pricing function SHA-256. It mainly consists of basic logical operators and relies solely on the computational speed of processor, therefore is perfectly suitable for multicore/conveyer implementation.

However, modern computers are not limited by the number of operations per second alone, but also by memory size. While some processors can be substantially faster than others, memory sizes are less likely to vary between machines.

Memory-bound price functions were first introduced by Abadi et al and were defined as "functions whose computation time is dominated by the time spent accessing memory." The main idea is to construct an algorithm allocating a large block of data ("scratchpad") within memory that can be accessed relatively slowly (for example, RAM) and "accessing an unpredictable sequence of locations" within it. A block should be large enough to make preserving the data more advantageous than recomputing it for each access. The algorithm also should prevent internal parallelism, hence N simultaneous threads should require N times more memory at once.

Dwork et al investigated and formalized this approach leading them to suggest another variant of the pricing function: "Mbound". One more work belongs to F. Coelho, who proposed the most effective solution: "Hokkaido".

To our knowledge the last work based on the idea of pseudo-random searches in a big array is the algorithm known as "scrypt" by C. Percival. Unlike the previous functions it focuses on key derivation, and not proof-of-work systems. Despite this fact scrypt can serve our purpose: it works well as a pricing function in the partial hash conversion problem such as SHA-256 in Bitcoin.
By now scrypt has already been applied in Litecoin and some other Bitcoin forks. However, its implementation is not really memory-bound: the ratio "memory access time / overall time" is not large enough because each instance uses only 128 KB. This permits GPU miners to be roughly 10 times more effective and continues to leave the possibility of creating relatively cheap but highly-efficient mining devices.
Moreover, the scrypt construction itself allows a linear trade-off between memory size and CPU speed due to the fact that every block in the scratchpad is derived only from the previous. For example, you can store every second block and recalculate the others in a lazy way, i.e. only when it becomes necessary.

Literature review

Our scheme relies on the cryptographic primitive called a group signature. First presented by D. Chaum and E. van Heyst, it allows a user to sign his message on behalf of the group. After signing the message the user provides (for verification purposes) not his own single public key, but the keys of all the users of his group. A verifier is convinced that the real signer is a member of the group, but cannot exclusively identify the signer.
The original protocol required a trusted third party (called the Group Manager), and he was the only one who could trace the signer. The next version called a ring signature, introduced by Rivest et al. in, was an autonomous scheme without Group Manager and anonymity revocation. Various modifications of this scheme appeared later: linkable ring signature allowed to determine if two signatures were produced by the same group member, traceable ring signature limited excessive anonymity by providing possibility to trace the signer of two messages with respect to the same metainformation (or "tag").
A similar cryptographic construction is also known as a ad-hoc group signature. It emphasizes the arbitrary group formation, whereas group/ring signature schemes rather imply a fixed set of members.
For the most part, our solution is based on the work "Traceable ring signature" by E. Fujisaki and K. Suzuki. In order to distinguish the original algorithm and our modification we will call the latter a one-time ring signature, stressing the user's capability to produce only one valid signature under his private key. We weakened the traceability property and kept the linkability only to provide one-timeness: the public key may appear in many foreign verifying sets and the private key can be used for generating a unique anonymous signature. In case of a double spend attempt these two signatures will be linked together, but revealing the signer is not necessary for our purposes.

Discussion

We have described the RPCA, which satisfies the conditions of correctness, agreement, and utility which we have outlined above. The result is that the zkFUND Protocol is able to process secure and reliable transactions in a matter of seconds: the length of time required for one round of consensus to complete. These transactions are provably secure up to the bounds outlined in section 3, which, while not the strongest available in the literature for Asynchronous Byzantine consensus, do allow for rapid convergence and flexibility in network membership. When taken together, these qualities allow the zkFUND Network to function as a fast and low-cost global payment network with well-understood security and reliability properties.
While we have shown that the zkFUND Protocol is provably secure so long as the bounds described in equations 1 and 3 are met, it is worth noting that these are maximal bounds, and in practice the network may be secure under significantly less stringent conditions. It is also important to recognize, however, that

satisfying these bounds is not inherent to the RPCA itself, but rather requires management of the UNLs of all users. The default UNL provided to all users is already sufficient, but should a user make changes to the UNL, it must be done with knowledge of the above bounds. In addition, some monitoring of the global network structure is required in order to ensure that the bound in equation 3 is met, and that agreement will always be satisfied.

We believe the RPCA represents a significant step forward for distributed payment systems, as the low-latency allows for many types of financial transactions previously made difficult or even impossible with other, higher latency consensus methods.

Questions and Intuition

Here are some questions that since these weeks, dreams asked me and I woke up sweating. But in fact it is OK.

```
Q. If you delete the transaction outputs, user cannot verify the
rangeproof and maybe a negative amount is created.

A. This is OK. For the entire transaction to validate all negative amounts
must have been destroyed. User have SPV security only that no illegal
inflation happened in the past, but the user knows that _at this time_ no
inflation occurred.
```

```
Q. If you delete the inputs, double spending can happen.

A. In fact, this means: maybe someone claims that some unspent output was
spent in the old days. But this is impossible, otherwise the sum of the
combined transaction could not be zero.

An exception is that if the outputs are amount zero, it is possible to
make two that are negatives of each other, and the pair can be revived
without anything breaks. So to prevent consensus problems, outputs
0-amount should be banned. Just add H at each output, now they all amount
to at least 1.
```

Future Research

Here are some questions I can not answer at the time of this writing.

```
1. What script support is possible? We would need to translate script
operations into some sort of discrete logarithm information.
```

```
2. We require user to check all k*G values, when in fact all that is
needed is that their sum is of the form k*G. Instead of using signatures
```

```
is there another proof of discrete logarithm that could be combined?
```

```
3. There is a denial-of-service option when a user downloads the chain,
the peer can give gigabytes of data and list the wrong unspent outputs.
The user will see that the result do not add up to 0, but cannot tell
where the problem is.

For now maybe the user should just download the blockchain from a Torrent
or something where the data is shared between many users and is reasonably
likely to be correct.
```

Conclusion

We have proposed a system for electronic transactions without relying on trust. We started with the usual framework of coins made from digital signatures, which provides strong control of ownership, but is incomplete without a way to prevent double-spending. To solve this, we proposed a peer-to-peer network using proof-of-work to record a public history of transactions that quickly becomes computationally impractical for an attacker to change if honest nodes control a majority of CPU power. The network is robust in its unstructured simplicity. Nodes work all at once with little coordination. They do not need to be identified, since messages are not routed to any particular place and only need to be delivered on a best effort basis. Nodes can leave and rejoin the network at will, accepting the proof-of-work chain as proof of what happened while they were gone. They vote with their CPU power, expressing their acceptance of valid blocks by working on extending them and rejecting invalid blocks by refusing to work on them. Any needed rules and incentives can be enforced with this consensus mechanism.

The zkFUND protocol was originally conceived as an upgraded version of a cryptocurrency, providing advanced features such as on-blockchain escrow, withdrawal limits, financial contracts, gambling markets and the like via a highly generalized programming language. The zkFUND protocol would not "support" any of the applications directly, but the existence of a Turing-complete programming language means that arbitrary contracts can theoretically be created for any transaction type or application. What is more interesting about zkFUND, however, is that the zkFUND protocol moves far beyond just currency. Protocols around decentralized file storage, decentralized computation and decentralized prediction markets, among dozens of other such concepts, have the potential to substantially increase the efficiency of the computational industry, and provide a massive boost to other peer-to-peer protocols by adding for the first time an economic layer. Finally, there is also a substantial array of applications that have nothing to do with money at all.

The concept of an arbitrary state transition function as implemented by the zkFUND protocol provides for a platform with unique potential; rather than being a closed-ended, single-purpose protocol intended for a specific array of applications in data storage, gambling or finance, zkFUND is open-ended by design, and we believe that it is extremely well-suited to serving as a foundational layer for a very large number of both financial and non-financial protocols in the years to come.

We have investigated the major flaws in Bitcoin and proposed some possible solutions. These ad- vantageous features and our ongoing development make new electronic cash system zkFUND a serious rival to Bitcoin, outclassing all its forks.

Nobel prize laureate Friedrich Hayek in his famous work proves that the existence of con- current independent

currencies has a huge positive effect. Each currency issuer (or developer in our case) is trying to attract users by improving his product. Currency is like a commodity: it can have unique benefits and shortcomings and the most convenient and trusted currency has the greatest demand. Suppose we had a currency excelling Bitcoin: it means that Bitcoin would develop faster and become better. The biggest support as an open source project would come from its own users, who are interested in it.

We do not consider zkFUND as a full replacement to Bitcoin. On the contrary, having two (or more) strong and convenient currencies is better than having just one. Running two and more different projects in parallel is the natural flow of electronic cash economics.

Creating a network of micropayment channels enables bitcoin scalability, micropayments down to the satoshi, and near-instant transactions. These channels represent real Bitcoin transactions, using the Bitcoin scripting opcodes to enable the transfer of funds without risk of counterparty theft, especially with long-term miner risk mitigations.

If all transactions using Bitcoin were on the blockchain, to enable 7 billion people to make two transactions per day, it would require 24GB blocks every ten minutes at best (presuming 250 bytes per transaction and 144 blocks per day). Conducting all global payment transactions on the blockchain today implies miners will need to do an incredible amount of computation, severely limiting bitcoin scalability and full nodes to a few centralized processors.

If all transactions using Bitcoin were conducted inside a network of micropayment channels, to enable 7 billion people to make two channels per year with unlimited transactions inside the channel, it would require 133 MB blocks (presuming 500 bytes per transaction and 52560 blocks per year). Current generation desktop computers will be able to run a full node with old blocks pruned out on 2TB of storage.

With a network of instantly confirmed micropayment channels whose payments are encumbered by timelocks and hashlock outputs, Bitcoin can scale to billions of users without custodial risk or blockchain centralization when transactions are conducted securely off-chain using bitcoin scripting, with enforcement of non-cooperation by broadcasting signed multisignature transactions on the blockchain.

References

W. Dai, "b-money," http://www.weidai.com/bmoney.txt, 1998.

H. Massias, X.S. Avila, and J.-J. Quisquater, "Design of a secure timestamping service with minimal trust requirements," In 20th Symposium on Information Theory in the Benelux, May 1999.

S. Haber, W.S. Stornetta, "How to time-stamp a digital document," In Journal of Cryptology, vol 3, no 2, pages 99-111, 1991.

D. Bayer, S. Haber, W.S. Stornetta, "Improving the efficiency and reliability of digital time-stamping," In Sequences II: Methods in Communication, Security and Computer Science, pages 329-334, 1993.

S. Haber, W.S. Stornetta, "Secure names for bit-strings," In Proceedings of the 4th ACM Conference on Computer and Communications Security, pages 28-35, April 1997.

A. Back, "Hashcash - a denial of service counter-measure," http://www.hashcash.org/papers/hashcash.pdf, 2002.

R.C. Merkle, "Protocols for public key cryptosystems," In Proc. 1980 Symposium on Security and Privacy, IEEE Computer Society, pages 122-133, April 1980.

W. Feller, "An introduction to probability theory and its applications," 1957.

Koinster, "White Paper Generator," http://whitepaper.koinster.com, 2017.

Nakamoto, Satoshi. "Bitcoin: A peer-to-peer elec- tronic cash system." Consulted 1.2012 (2008): 28.

Lamport, Leslie, Robert Shostak, and Marshall Pease. "The Byzantine generals problem." ACM Transactions on Programming Languages and Systems (TOPLAS) 4.3 (1982): 382-401.

Attiya, C., D. Dolev, and J. Gill. "Asynchronous Byzantine Agreement." Proc. 3rd. Annual ACM Symposium on Principles of Distributed Computing. 1984.

Fischer, Michael J., Nancy A. Lynch, and Michael S. Paterson. "Impossibility of distributed consensus with one faulty process." Journal of the ACM (JACM) 32.2 (1985): 374-382.

Martin, J-P., and Lorenzo Alvisi. "Fast byzantine consensus." Dependable and Secure Computing, IEEE Transactions on 3.3 (2006): 202-215.

Alchieri, Eduardo AP, et al. "Byzantine consensus with unknown participants." Principles of Distributed Systems. Springer Berlin Heidelberg, 2008. 22-40.

Intrinsic value: http://bitcoinmagazine.com/8640/an-exploration-of-intrinsic-value-what-it-is-why-bitcoin-doesnt-have-it-and-why-bitcoin-does-have-it/

Smart property: https://en.bitcoin.it/wiki/Smart_Property

Smart contracts: https://en.bitcoin.it/wiki/Contracts

B-money: http://www.weidai.com/bmoney.txt

Reusable proofs of work: http://www.finney.org/~hal/rpow/

Secure property titles with owner authority: http://szabo.best.vwh.net/securetitle.html

Bitcoin whitepaper: http://bitcoin.org/bitcoin.pdf

Namecoin: https://namecoin.org/

Zooko's triangle: http://en.wikipedia.org/wiki/Zooko's_triangle

Colored coins whitepaper: https://docs.google.com/a/buterin.com/document/d/1AnkP_cVZTCMLIzw4DvsW6M8Q2JC0lIzrTLuoWu2z1BE/edit

Mastercoin whitepaper: https://github.com/mastercoin-MSC/spec

Decentralized autonomous corporations, Bitcoin Magazine: http://bitcoinmagazine.com/7050/bootstrapping-a-decentralized-autonomous-corporation-part-i/

Simplified payment verification: https://en.bitcoin.it/wiki/Scalability#Simplifiedpaymentverification

Merkle trees: http://en.wikipedia.org/wiki/Merkle_tree

Patricia trees: http://en.wikipedia.org/wiki/Patricia_tree

GHOST: http://www.cs.huji.ac.il/~avivz/pubs/13/btc_scalability_full.pdf

StorJ and Autonomous Agents, Jeff Garzik: http://garzikrants.blogspot.ca/2013/01/storj-and-bitcoin-autonomous-agents.html

Mike Hearn on Smart Property at Turing Festival: http://www.youtube.com/watch?v=Pu4PAMFPo5Y

Ethereum RLP: https://github.com/ethereum/wiki/wiki/%5BEnglish%5D-RLP

Ethereum Merkle Patricia trees: https://github.com/ethereum/wiki/wiki/%5BEnglish%5D-Patricia-Tree

Peter Todd on Merkle sum trees: http://sourceforge.net/p/bitcoin/mailman/message/31709140

http://bitcoin.org.

https://en.bitcoin.it/wiki/Category:Mixing Services.

http://blog.ezyang.com/2012/07/secure-multiparty-bitcoin-anonymization.

https://bitcointalk.org/index.php?topic=279249.0.

http://msrvideo.vo.msecnd.net/rmcvideos/192058/dl/192058.pdf.

https://github.com/bitcoin/bips/blob/master/bip-0034.mediawiki#Specification.

https://github.com/bitcoin/bips/blob/master/bip-0016.mediawiki#Backwards Compatibility.

https://en.bitcoin.it/wiki/Mining hardware comparison.

https://github.com/bitcoin/bips/blob/master/bip-0050.mediawiki.

http://luke.dashjr.org/programs/bitcoin/files/charts/branches.html.

https://bitcointalk.org/index.php?topic=196259.0.

https://en.bitcoin.it/wiki/Contracts.

https://en.bitcoin.it/wiki/Script.

http://litecoin.org.

Martin Abadi, Michael Burrows, and Ted Wobber. Moderately hard, memory-bound functions. In NDSS, 2003.

Ben Adida, Susan Hohenberger, and Ronald L. Rivest. Ad-hoc-group signatures from hijacked keypairs. In in DIMACS Workshop on Theft in E-Commerce, 2005.

Man Ho Au, Sherman S. M. Chow, Willy Susilo, and Patrick P. Tsang. Short linkable ring signatures revisited. In EuroPKI, pages 101-115, 2006.

Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. J. Cryptographic Engineering, 2(2):77-89, 2012.

David Chaum and Eug`ene van Heyst. Group signatures. In EUROCRYPT, pages 257-265, 1991.

Fabien Coelho. Exponential memory-bound functions for proof of work protocols. IACR Cryptology ePrint Archive, 2005:356, 2005.

Ronald Cramer, Ivan Damg ÌŠard, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In CRYPTO, pages 174-187, 1994.

Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. In CRYPTO, pages 426-444, 2003.

Eiichiro Fujisaki. Sub-linear size traceable ring signatures without random oracles. In CT- RSA, pages 393-415, 2011.

Eiichiro Fujisaki and Koutarou Suzuki. Traceable ring signature. In Public Key Cryptogra- phy, pages 181-200, 2007.

Jezz Garzik. Peer review of "quantitative analysis of the full bitcoin transaction graph". https://gist.github.com/3901921, 2012.

Joseph K. Liu, Victor K. Wei, and Duncan S. Wong. Linkable spontaneous anonymous group signature for ad hoc groups (extended abstract). In ACISP, pages 325-335, 2004.

Joseph K. Liu and Duncan S. Wong. Linkable ring signatures: Security models and new schemes. In ICCSA (2), pages 614-623, 2005.

Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In IEEE Symposium on Security and Privacy, pages 397- 411, 2013.

Micha Ober, Stefan Katzenbeisser, and Kay Hamacher. Structure and anonymity of the bitcoin transaction graph. Future internet, 5(2):237-250, 2013.

Tatsuaki Okamoto and Kazuo Ohta. Universal electronic cash. In CRYPTO, pages 324-337, 1991.

Marc Santamaria Ortega. The bitcoin transaction graph â€" anonymity. Masterâ€™s thesis, Universitat Oberta de Catalunya, June 2013.

Colin Percival. Stronger key derivation via sequential memory-hard functions. Presented at BSDCanâ€™09, May 2009.

Fergal Reid and Martin Harrigan. An analysis of anonymity in the bitcoin system. CoRR, abs/1107.4524, 2011.

Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In ASIACRYPT, pages 552-565, 2001.

Dorit Ron and Adi Shamir. Quantitative analysis of the full bitcoin transaction graph. IACR Cryptology ePrint Archive, 2012:584, 2012.

Meni Rosenfeld. Analysis of hashrate-based double-spending. 2012.

Maciej Ulas. Rational points on certain hyperelliptic curves over finite fields. Bulletin of the Polish Academy of Sciences. Mathematics, 55(2):97-104, 2007.

Qianhong Wu, Willy Susilo, Yi Mu, and Fangguo Zhang. Ad hoc group signatures. In IWSEC, pages 120-135, 2006.

https://people.xiph.org/~greg/confidential_values.txt

https://bitcointalk.org/index.php?topic=279249.0

https://cryptonote.org/whitepaper.pdf

https://eprint.iacr.org/2015/1098.pdf

https://download.wpsoftware.net/bitcoin/wizardry/horasyuanmouton-owas.pdf

http://blockstream.com/sidechains.pdf

http://fr.harrypotter.wikia.com/wiki/Sortilege_de_Langue_de_Plomb

https://bitcointalk.org/index.php?topic=281848.0

Satoshi Nakamoto. Bitcoin: A Peer-to-peer Electronic Cash System. https://bitcoin.org/bitcoin.pdf, Oct 2008.

Manny Trillo. Stress Test Prepares VisaNet for The Most Wonderful Time of The Year http://www.visa.com/blogarchives/us/2013/10/10/stress-test-prepares-visanet-for-the-most-wonderful-time-of-the-year/index.html, Oct 2013.

Bitcoin Wiki. Contracts: Example 7: Rapidly-adjusted (micro)payments to a pre-determined party. https://en.bitcoin.it/wiki/Contracts#Example_7:_Rapidly-adjusted_.28micro.29payments_to_a_pre-determined_party.

bitcoinj. Working with micropayment channels. https://bitcoinj.github.io/working-with-micropayments.

Leslie Lamport. The Part-Time Parliament. ACM Transactions on Computer Systems, 21(2):133â€“169, May 1998.

Leslie Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. Communications of the ACM, 21(7):558â€“565, Jul 1978.

Alex Akselrod. Draft. https://en.bitcoin.it/wiki/User:Aakselrod/Draft, Mar 2013.

Alex Akselrod. ESCHATON. https://gist.github.com/aakselrod/9964667, Apr 2014.

Peter Todd. Near-zero fee transactions with hub-and-spoke micropayments. http://sourceforge.net/p/bitcoin/mailman/message/33144746/, Dec 2014.

C.J. Plooy. Combining Bitcoin and the Ripple to create a fast, scalable, decentralized, anonymous, low-trust payment network. http://www.ultimatestunts.nl/bitcoin/ripple_bitcoin_draft_2.pdf, Jan 2013.

BitPay. Impulse. http://impulse.is/impulse.pdf, Jan 2015.

Mark Friedenbach. BIP 0068: Consensus-enforced transaction replacement signaled via sequence numbers (relative locktime). https://github.com/bitcoin/bips/blob/master/bip-0068.mediawiki, May 2015.

Mark Friedenbach BtcDrak and Eric Lombrozo. BIP 0112: CHECKSEQUENCEVERIFY. https://github.com/bitcoin/bips/blob/master/bip-0112.mediawiki, Aug 2015.

Jonas Schnelli. What does OP CHECKSEQUENCEVERIFY do? http://bitcoin.stackexchange.com/a/38846, Jul 2015.

Greg Maxwell (nullc). reddit. https://www.reddit.com/r/Bitcoin/comments/37fxqd/it_looks_like_blockstream_is_working_on_the/crmr5p2, May 2015.

Gavin Andresen. BIP 0016: Pay to Script Hash. https://github.com/bitcoin/bips/blob/master/bip-0016.mediawiki, Jan 2012.

Pieter Wuille. BIP 0032: Hierarchical Deterministic Wallets. https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki, Feb 2012.

Ilja Gerhardt and Timo Hanke. Homomorphic Payment Addresses and the Pay-to-Contract Protocol. http://arxiv.org/abs/1212.3257, Dec 2012.

Nick Szabo. Formalizing and Securing Relationships on Public Networks. http://szabo.best.vwh.net/formalize.html, Sep 1997.

Donate

Feel free to support this project by donating Bitcoin here:

`1MUnEgxuMykj6z6rKnpozm8rCWpqdpSjen`