Open Archive Toulouse Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of some Toulouse researchers and makes it freely available over the web where possible.

This is an author's version published in: https://oatao.univ-toulouse.fr/20963

To cite this version :

Any correspondence concerning this service should be sent to the repository administrator: tech-oatao@listes-diff.inp-toulouse.fr

Erasure code-based low storage blockchain node

Doriane Perard, Jérôme Lacan, Yann Bachy and Jonathan Detchart
ISAE-Supaero, Université de Toulouse, France
firstname.name@isae-supaero.fr

Abstract—The concept of a decentralized ledger usually implies that each node of a blockchain network stores the entire blockchain. However, in the case of popular blockchains, which each weigh several hundreds of GB, the large amount of data to be stored can incite new or low-capacity nodes to run *lightweight* clients. Such nodes do not participate to the global storage effort and can result in a centralization of the blockchain by very few nodes, which is contrary to the basic concepts of a blockchain.

To avoid this problem, we propose new *low storage* nodes that store a reduced amount of data generated from the blockchain by using erasure codes. The properties of this technique ensure that any block of the chain can be easily rebuilt from a small number of such nodes. This system should encourage low storage nodes to contribute to the storage of the blockchain and to maintain decentralization despite of a globally increasing size of the blockchain. This system paves the way to new types of blockchains which would only be managed by low capacity nodes.

I. INTRODUCTION

One of the most interesting properties of blockchains is their decentralized nature allowing to avoid central authorities. Classically, each node participating to the blockchain must maintain it by participating to the consensus when inserting a new block and by storing the entire blockchain. However, the success of blockchains such as Bitcoin or Ethereum has highlighted a scalability problem, making them less and less decentralized [1]. Indeed, the increasing size of these blockchains implies important efforts by medium capacity nodes:

1) CPU: each transaction has to be processed by every node.
2) Storage: the size of many popular blockchains is considerably increasing, which implicates an important storage effort for classic nodes.
3) Network load: Points 1) and 2) can lead to a reduced number of nodes storing the entire blockchain. This could increase the network load for the remaining full nodes.

Several solutions were recently proposed to cope with the general scalability problem of blockchains. For example, the Ethereum team [2] proposes a mechanism allowing to increase the blockchain's throughput. This solution consequently reduces the amount of nodes verifying each transaction without impacting the global security level. Doing so allows to process more transactions in parallel. Other solutions use auxiliary channels to debottleneck the network and save storage space, like Plasma Cash or lightning network [3], [4].

Considering the storage scalability problem, the most frequent solution consists in nodes running a lightweight client, thereby reducing their required storage effort. A side effect of this solution is an important network and processing load on full nodes satisfying light client requests.

To the best of our knowledge, blockchains only use data replication. The aim of this paper is to propose the use of erasure codes in order to store the blockchain. To do so, we introduce a new type of node, an erasure code-based low storage node, which stores coded fragments for each block of the blockchain. The main benefit is to allow a reduction of the storage effort required by each node. This should encourage nodes with low storage capacity to participate to the storage effort of the blockchain and thus result in two other benefits which are 1) the contribution to maintaining the decentralized property of the blockchain and 2) the reduction of network load on each node by multiplying the number of nodes.

This paper is structured as follows. First, Section II describes some related works concerning blockchain storage mechanisms and erasure codes for distributed storage systems. Then, Sections III and IV respectively present a high level and low-level description of erasure-code based low storage nodes, the main contribution of this paper. After, Sections V and VI present the interest of our low storage blockchain node and an analysis of the available parameters of our system. Finally, Section VII concludes this paper and proposes some perspectives.

II. RELATED WORKS

Let us consider blockchains as a sequence of variable sized blocks, virtually connected by their respective hashes and stored on a peer-to-peer network composed of several types of nodes. The first type, called *full nodes*, stores the entire blockchain.The second type, called *lightweight nodes*, only stores the most recent blocks and the hashes of every block of the chain. This type of node allows to participate to a blockchain without storing the entire chain when the size of the blockchain becomes important. In the rest of this section we will discuss three solutions allowing to face important data storage.

A. Replicated systems

A first solution allowing to reduce the required storing capacity of every node consists in storing blocks among

the nodes, in such a way that every block continues to be sustained by several nodes but not by all.

If one of the nodes storing a block becomes unavailable, the block continues to be stored by others. However, the global availability of this block is drastically reduced. In this situation, malicious people could be tempted to target parts of the blockchain that are poorly replicated.

B. Light nodes

A second solution to storage limitation is offered by light clients. Many blockchains define their own type of light client [5]–[8]. Light clients do not store the entire blockchain, but only parts of it, generally the most recent blocks. As an example we can cite Simplified Payment Verification (SPV) which was originally designed by bitcoin's creator in his whitepaper [9]. SPV allows nodes with limited resources to participate to a blockchain and perform payment verification, while only storing block headers. The node can not verify transactions by itself, but it requests block-content when necessary and uses the merkle tree to confirm the integrity of the data it received.

Nodes running light clients present many problems:

- They don't store every block, so availability is not improved. If all nodes would run such clients, the overall blockchain availability would be endangered.
- They never verify the entire blockchain, but instead rely their trust on other full nodes.

C. Erasure code and distributed storage systems

A third solution to data storage is inspired from classic distributed storage systems. Such systems often use erasure codes, like in RAID (Redundant Array of Independent Disks) [10], Clouds [11] or P2P systems [12]. The main idea is to encode the data into coded chunks which are distributed over several nodes of the system. The Storj system [13] uses a similar concept by storing coded and encrypted data in a P2P network that can be connected to a blockchain. The Swarm project [14] also uses erasure codes in a blockchain context, it's purpose is to provide a decentralized and redundant store for dapp code and data, on the ethereum blockchain.

In all these systems, erasure coding improves the trade-off between data availability and the amount of stored data.

To describe the concept behind erasure codes, let us first consider a replicated system, represented in Fig. 1, which stores 3 data blocks b_1, b_2 and b_3, these blocks are simply replicated for redundancy. Here, if a block and its replicated block are deleted or at least unavailable at a given time, this block can no longer be recovered.

When using erasure codes, the 3 redundancy blocks store linear combinations of the data blocks. With the same block unavailability, as shown in Fig. 2, we can observe that the 3 blocks can be recovered from the 3 available blocks by performing the inverse linear operations (we assume that the linear system is invertible).
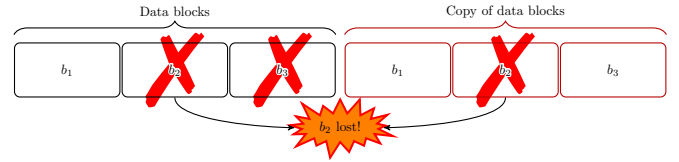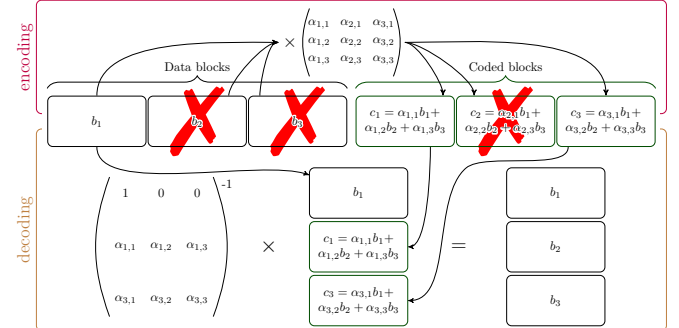

Fig. 1: Traditional system

Fig. 2: Simple example of coded storage system

Let us conclude that the use of erasure codes provides a better availability than classic data replication with the same amount of storage. The objective of this paper is to improve blockchain storage scalability by using erasure codes.

III. Low storage nodes

In this paper we introduce a new type of node, which is called *Low Storage (LS) node*. They are described in the next paragraph. The way they recover a given block and the way they interact with the blockchain network is presented in the following paragraphs.

A. Principle

The main contribution of LS nodes to the blockchain system is to store only some *coded fragments* of each block. These fragments are obtained by first splitting a block into fixed size fragments and then generating linear combinations of these fragments. Note that the combinations are specific to each block and to each node. This means that when several nodes store coded fragments from the same block, these will necessarily be different.

A detailed description of these operations is presented in Section IV. The number of coded fragments stored by an LS node can vary according to different factors such as the age of the block or the storage capabilities of the node.

B. Block recovery

When an LS node wants to recover a block from coded fragments stored by different LS nodes, it shall download coded fragments corresponding to the desired block upon several LS nodes. When the amount of coded fragments

downloaded is more important than the amount of fragments the block was initially split into, the node can access the block by performing the inverse linear combinations.

C. Integration of a new LS node in the network

When an LS node wants to join the network, it needs to download all the blocks of the blockchain. This can be done from one or more trusted full nodes or from several LS nodes. For each block, it verifies its validity according to the previous block and generates its own coded fragments. Then, it removes the complete block and only stores the hash of the block and the coded fragments. Since the operations done on a block are independent from the other blocks, it can process the entire blockchain in sequence. This has the advantage of not needing to store the entire blockchain at once, therefore LS nodes are able to verify the entire blockchain with a reduced storage capacity.

IV. CODING THE BLOCKCHAIN

In this section we will present the technical aspects of coding block fragments. Firstly, we will discuss how we apply erasure coding to our system. The coding system in our system follows the same principle as coding mechanisms used in distributed storage systems (cf. Fig. 2). Secondly we will present how a node can recover a block based on coded fragments.

Before describing the coding operation, let us first define some notations. We denote by $N^{(i)}$ the nodes of the network, where i is an unique identifier characterizing each node. We denote by $B^{(j)}$ the j^{th} block of the blockchain. We consider that the first block is $B^{(0)}$.

Let us denote by s_B the maximum size of a block of the blockchain. Let us define two integers k and r respectively corresponding to the number of fragments of a block and the number of coded fragments stored by a node. The choice of the values of k and r will be discussed in Section VI.

A. Coding the data

When a node wants to code a block $B^{(j)}$, it must first split the block in k fragments before applying a linear operation in order to obtain r coded fragments.

To generate r coded fragments from block $B^{(j)}$, the node $N^{(i)}$ first splits the block into k fragments $F_0^{(j)}, F_1^{(j)}, \ldots, F_{k-1}^{(j)}$. If necessary, before splitting the block, it is padded with zeros up to s_B in order to have a fixed fragment size s_B/k. Then, the node $N^{(i)}$ initializes a pseudo-random number generator with a seed defined as the concatenation of the binary expressions of i and j and generates the pseudo-random sequence $\{\alpha_0^{(i,j)}, \ldots, \alpha_{k \cdot r-1}^{(i,j)}\}$ of $k \cdot r$ coefficients in the finite field $\mathbb{F}_{2^m}$. We propose to set $m := 8$ because the field $\mathbb{F}_{2^8}$ is generally considered as a good tradeoff between the complexity of its operations and the diversity it provides. This choice will be explained in Section VI. In practical, the 256 elements of this finite field are associated by a bijection to the 256 bytes.

To build the coded fragment $\mathfrak{F}_u^{(i,j)}$, where $0 \leq u \leq r-1$, the node considers the k coefficients $\{\alpha_{k \cdot u}^{(i,j)}, \ldots, \alpha_{(k+1) \cdot u-1}^{(i,j)}\}$ and performs the following linear combination:

$$\mathfrak{F}_u^{(i,j)} = \alpha_{k \cdot u}^{(i,j)} \cdot F_0^{(j)} + \ldots, \alpha_{(k+1) \cdot u-1}^{(i,j)} \cdot F_{k-1}^{(j)}$$

From a more practical point of view, if $f_{l,v}^{(j)}$ denotes the finite field element associated to the v^{th} byte of the fragment $B_l^{(j)}$, then the v^{th} byte of the coded fragment $\mathfrak{F}_u^{(i,j)}$ is defined by the finite field element $\mathfrak{f}_{u,v}^{(i,j)}$ computed as follows:

$$\mathfrak{f}_{u,v}^{(i,j)} = \alpha_{k \cdot u}^{(i,j)} \cdot f_{0,v}^{(j)} + \ldots, + \alpha_{(k+1) \cdot u-1}^{(i,j)} \cdot f_{k-1,v}^{(j)}$$

A more detailed presentation of the linear combinations of byte packets in finite fields is presented in [15].

Finally, the node removes the block $B^{(j)}$ and replaces it by the coded fragments $\mathfrak{F}_0^{(i,j)}, \ldots, \mathfrak{F}_{r-1}^{(i,j)}$.

B. Recovering the data

Practically, a node $N^{(i)}$ downloads coded fragments $\mathfrak{F}_0^{(i_0,j)}, \ldots, \mathfrak{F}_{r-1}^{(i_{r-1},j)}$ from several nodes $N^{(i_0)}, \ldots, N^{(i_{r-1})}$. Note that it is possible to request multiple coded fragments from the same node. Since the node $N^{(i)}$ knows the values of $i_0, \ldots, i_{r-1}$ and j, it is able to initialize the correct pseudo random number generator with the binary expressions of $i_0, \ldots, i_{r-1}$ and j and thus can regenerate the coefficients α used when $N^{(i_0)}, \ldots, N^{(i_{r-1})}$ built their coded fragments. So, $N^{(i)}$ knows the r linear equations used to build the r coded fragments from the initial fragments. After downloading a sufficient number of coded fragments, the node will have enough equations to invert the linear system and recover the k fragments (and thus the block) from the coded fragments.

In Section VI we will show that, with the field $\mathbb{F}_{2^8}$, in most cases the linear system can be inverted as soon as k coded fragments are downloaded.

V. LOW STORAGE NODE INTEREST

The main objective of our system is to allow any node to contribute to an entire blockchain with a reduced storage effort. In this section we will discuss the different interests of our system compared to traditional blockchain nodes. Firstly we will compare the storage effort required by traditional vs coded blockchains. We will show how low storage blockchain nodes can relieve full nodes, and improve trust in the network. Secondly we will show how our system can improve the information security in terms of availability and integrity.

A. Storage effort scalability

Let c be the compression factor, defined by $c := k/r$. Note that a node that stores r coded fragments must download $k - r$ extra coded fragments in order to recover a block. So, the higher the compression factor, the more coded fragments the node must download to be able to

decode the entire block. The choice of these parameters will be discussed in section VI.

When $r = k$, $c = 1$ which means that the storage effort is not reduced. In this case, the data can be rebuilt without downloading additional coded fragments. When r decreases down to 1, the compression factor increases up to k. Considering that each node must store at least one coded fragment per block, the maximum compression factor for each block is k.

One of the interests of our system is its scalability. Indeed, each node can adapt r according to, for example, the age of a block by simply removing some of its stored coded fragments without re-calculating them.

B. Availability

One of the main goals of our system is to improve the global availability and sustainability of a blockchain. In order to achieve this, our system reduces the storage effort required by a single node while participating to the global storage effort of the entire blockchain.

With our system we consider almost any node stores at least one coded fragment of every block. Thereby, even if each block is not stored in its initial form, every block is entirely stored among the different nodes, each one storing different fragments. In most cases, the initial block can be recovered by combining any k coded fragments provided by any node of the network. For a system with a large amount of nodes, any node can leave the system or be unreachable without significantly impacting the availability.

To evaluate the availability, let us denote n the number of nodes. Since each node can store a variable number of fragments, we note r_i the number of fragments stored by i^{th} node. We assume that the r_i can be modeled as realizations of an independent discrete random variable following a probability distribution f. Since a block can be recovered when at least k coded fragments are available, the probability $\overline{P_c}$ of block irrecoverability from coded fragments is :

$$\overline{P_c} = p\left(\sum_{i=0}^{n-1} r_i < k\right) = \sum_{u=0}^{k-1} p\left(\sum_{i=0}^{n-1} r_i = u\right) = \sum_{u=0}^{k-1} f^{*^n}(u) \quad (1)$$

where $f^{*^n} = \overbrace{f * f * \cdots * f}^{n \text{ times}}$ and $*$ the convolution between two distribution probabilities.

To compare with a system based on random replication nodes with the same level of storage, we assume that each node stores each block with a probability of $1/c$. The probability that a block is *not* stored by a node is thus $1 - 1/c$. By assuming that each node independently chooses which block to store, the probability $\overline{P_r}$ that a block is stored by no node at all is equal to:
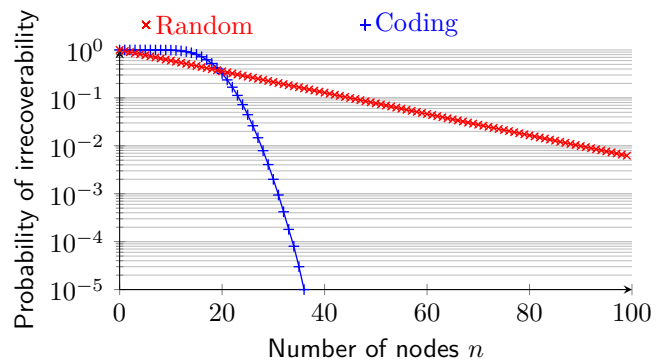
$$\overline{P_r} = (1 - \frac{1}{c})^n \quad (2)$$

Fig. 3: Probability one chosen block is irrecoverable, depending on n, with $k = 100$

Let us compare these two systems with a numerical example. Let us assume that, in the coded system, the number r of coded fragments stored by each node follows a geometric distribution $f_p(r) = p \cdot (1 - p)^{r-1}$. We then have:

$$f_p^{*^n}(u) = \frac{(u-1)!}{(n-1)!(u-n)!} \cdot p^n \cdot (1-p)^{u-n} \quad (3)$$

In our example, we assume that $k = 100$ and that the average compression factor of the block is $c = 20$. So, on average, each node stores $r = 5$ coded fragments. Since the mean of a geometric distribution of parameter p is $1/p$, we have $p = 0.2$.

With these parameters, Fig. 3 shows the probability of block irrecoverability from coded fragments for the two systems, by using Eqs. 1 and 3 for the coded system and Eq. 2 for the random replication system.

With the coded system and this probability law, starting from 37 nodes the probability to access a block equals 1 (with 5 digits precision), but with the randomly replicated it is only from 238 nodes.

Regardless of the real probability law, with low storage nodes every node has to store at least $r = 1$ fragments, so with only k nodes we are sure to access the block.

Moreover, if we consider every node has the same probability of failure, using our mechanism, the overall probability of blockchain unavailability will drop. Indeed, considering the network storage capacity remains identical, our method distributes data over many more nodes, thereby distributing and reducing the global failure rate of the blockchain.

C. Integrity

This section will discuss how our system not only conserves the integrity of the blockchain but can improve trust by making data tampering more difficult.

1) Hash-collision with coded fragments: By construction, the integrity of a blockchain is guaranteed by the hash contained in each block. This hash validates the previous block of the chain and mechanically the entire

blockchain. Therefore, any malicious node wanting to tamper a block of the blockchain would need to find a hash-collision in order to conserve the same hash result for the tampered block.

Since the malicious node can only tamper some of the coded fragments used for the decoding process, he must figure out which correct fragments will be used for the decoding process. Then, he must search a hash-collision such that the result of the decoding process of the correct and the tampered fragments has the same hash as the initial block. This problem is at least as difficult as the classical hash-collision problem. So the use of erasure codes does not only reduce but rather increases the integrity of the system.

2) Malicious node identification: Erasure codes also provide solutions to identify malicious nodes tampering blocks. Tampered blocks will not be compatible and coherent with any set of k valid coded fragments. This property allows each node to check the validity of other nodes and thereby improves the global security of the system. Moreover, this would allow the use of fraud proofs to report malicious nodes like detailed in [16].

3) Blockchain verification: As explained in IV-A, the seed used by the pseudo-random number generator to generate the coefficients of the linear combination depends on the identity of the node. When decoding a block, each node uses the seeds belonging to the nodes from which the coded fragments were downloaded. When storing coded fragments of a block, these are necessarily re-coded with the local seed.

When a node wants to download coded fragments and decode the block, it uses i, the unique node identifier to regenerate the coefficients.

The only way to obtain coded fragments coded with the correct coefficients, is to entirely code them from plain blocks. If a node does not follow this principle, the regenerated coefficients won't match and the decoding process will fail. This constraint forces each low storage node to verify the entire blockchain, which in turn will improve the overall integrity.

D. Network load balancing

We expect low storage nodes will contribute to a gain of the number of nodes, thereby improving the distribution of the blockchain over it's network.

As an example let us consider 1000 light nodes each connected with a 10Mbps network interface and only 5 full nodes serving the entire blockchain. In this situation each full node would have to be able to serve blocks at:

$$\frac{10Mbps \cdot 1000}{5} = 2Gbps$$

in order to satisfy the 1000 light nodes. With our method, if we consider each node stores 5 ($k = 100$, $c = 20$) coded fragments, with the same storage effort allowing 5 full nodes, our network can provide 100 ($5 * c = 5 * 20$) low storage nodes. In this situation each node would initiate 20

connections at 0.5Mbps (10Mbps / 20) which totals 20000 connections for all the nodes and therefore each full node would have to be able to serve blocks at only:

$$\frac{\frac{10Mbps}{20} \cdot 20000}{1000} = 100Mbps$$

in order to satisfy the 1000 nodes. This simple example clearly illustrates how our low storage nodes can decongest the network. Naturally connection establishment time and other overhead must also be considered when realizing a full network load study.

VI. ANALYSIS OF THE PARAMETERS

On each node, each block, and thus k fragments, are replaced by r coded fragments. One of the challenges is to determine k and r, with the best compromise between compression and complexity. In this section, we will present some consequences when varying these parameters. Every blockchain and every type of node is different, it is therefore impossible to determine a generic value for k and r but a compromise must be found according to the system's constraints. Before discussing fragmentation and compression factor we will consider finite field size and processing complexity, since these two aspects will influence the choice of k and r.

A. Size of the finite field

The choice of the finite field used to perform linear combinations mainly impacts two parameters: the probability of block recovery from downloaded coded fragments (which is better with a large finite field) and the encoding and decoding complexities (which is smaller with a small finite field).

Reference [17] shows that when the coefficients of the linear combinations are chosen in the field $\mathbb{F}_{2^m}$, the probability of recovering a block from $k + s$ downloaded fragments can be approximated by $1 - 2^{-m \cdot (s+1)}$.

Classically, for practical implementations, the possible sizes of the finite field q are 2^{2^i}, i.e. 2, 4, 16, 256, 65536. Since extremely fast implementations can be developed for $2^m = 256$ [18], we propose to use this value in our system.

B. Processing complexity

The complexity of encoding consists in multiplying a $k \cdot r$-matrix by the k original fragments. Then, there is $k \cdot r \cdot s_B / k = r \cdot s_B$ operations. So, the encoding complexity does not depend on k, but only on r.

For decoding, the complexity consists in inverting a $k \cdot k$-matrix, and then multiplying it by the k coded fragments. The pseudo-random matrix inversion has a complexity in $O(k^3)$. So the number of operations is $O(k^3) + k^2 \cdot s_B / k = O(k^3) + k \cdot s_B$ and thus depends only on k. If the size of the block is large compared to k, then the matrix-vector multiplication ($k \cdot$ block size) is the most complex operation.

Fig. 4 shows the encoding and decoding speeds of 1MB blocks, with different values of k and r. We used the
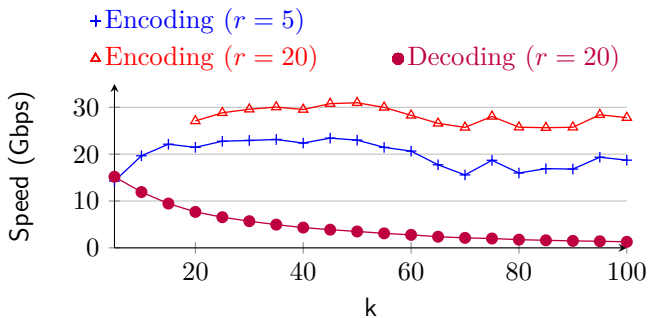
Fig. 4: Encoding and decoding speed, with block size = 1Mb

erasure code implementation of [18] with avx2 instructions on a core i7-6700@3.2 GHz CPU. This graph allows to conclude that the processing cost is negligible. Indeed, coding speed is always greater than 15 Gbps for $r \leq 20$. Decoding speed is lower, but greater than 3.5 Gbps for $k \leq 50$ and equals to 1.5 Gbps for $k = 100$.

C. Fragmentation

As described in Section IV-A, each node, in order to generate its r coded fragments, will split the initial block into k fragments. The choice of this parameter can be different for each blockchain, but it must be the same for every user of the same blockchain.

Overall we can conclude that k should be chosen as big as possible, without exceeding $\frac{2^m}{2} = 128$ ($2^m = 256$, size of our finite field). If s_B, the maximum size of a block of the blockchain, is too small, it possible to group a small amount of blocks before coding them. Doing so will allow to increase k and take a better benefit of low storage nodes.

D. Compression factor

The compression factor $c = k/r$ indicates by how much the storage effort of a node adopting the low storage mechanism can be reduced. Considering k is system-wide determined, c can only be chosen by varying r.

The choice of r is up to the end user and will depend on the type of user. Choosing r high will improve block recovery and reduce network load, also will this improve the overall blockchain availability as increasing r increases the storage effort of a node. Choosing r low, will reduce coding complexity and compression factor. Globally r must be chosen according to the node's capacities.

VII. CONCLUSION

The increasing popularity of blockchains reveals all the possibilities a decentralized ledger can offer. Nonetheless blockchains also have their limits and scalability rapidly appears problematic.

In this paper, we introduce a new type of blockchain node, called Low Storage node, with the aim of proposing an alternative between light nodes and full nodes. This type of node uses erasure codes in order to allow users with reduced storage capacity to contribute to the blockchain, without storing the entire blockchain, but instead only linear combinated fragments of each block. Eventhough blocks are no longer stored in their original form, they can still be recovered by downloading fragments from other nodes and performing the inverse linear operation.

We have shown that our system can improve the overall availability of a blockchain without impacting data integrity. We believe our system will encourage new nodes to participate to the network, thereby relieving existing nodes. Moreover, CPU ressources required for the coding process are negligible when coding at the same speed as average home internet connections allow to download the blockchain.

As future work we plan to explore an IoT compatible blockchain model. We also plan to study how to exempt nodes from verifying the entire blockchain when joining the network.

REFERENCES

[1] A. E. Gencer, S. Basu, I. Eyal, R. van Renesse, and E. G. Sirer, "Decentralization in bitcoin and ethereum networks," 2018. [Online]. Available: arxiv.org/abs/1801.03998
[2] V. Buterin, "On sharding blockchains," github.com/ethereum/wiki/wiki/Sharding-FAQ, November 2017.
[3] J. Poon and T. Dryja, "The bitcoin lightning network: Scalable off-chain instant payments," lightning.network/lightning-network-paper.pdf, January 2016.
[4] J. Poon and V. Buterin, "Plasma: Scalable autonomous smart contracts," plasma.io/plasma.pdf, August 2017.
[5] J. McKinney, "Light client protocol," github.com/ethereum/wiki/wiki/Light-client-protocol, November 2017.
[6] P. developer team, "Getting synced," github.com/paritytech/parity/wiki/Getting-Synced, January 2018.
[7] E. developer team, "Electrum," electrum.org, 2017.
[8] B. core developers, "Bitcoin core version 0.12.0 released," bitcoin.org/en/release/v0.12.0#wallet-pruning, February 2016.
[9] S. Nakamoto, bitcoin.org/bitcoin.pdf, 2009.
[10] J. S. Plank, "A tutorial on reed-solomon coding for fault-tolerance in raid-like systems," *Softw. Pract. Exper.*, vol. 27, no. 9, pp. 995–1012, Sep. 1997.
[11] J. Li and B. Li, "Erasure coding for cloud storage systems: A survey," *Tsinghua Science and Technology*, vol. 18, no. 3, pp. 259–272, June 2013.
[12] H. Weatherspoon and J. Kubiatowicz, "Erasure coding vs. replication: A quantitative comparison," in *IPTPS*, ser. Lecture Notes in Computer Science, vol. 2429. Springer, 2002, pp. 328–338.
[13] S. Wilkinson and al., "Storj: A peer-to-peer storage network," december 2014.
[14] V. Trón, A. Fischer, D. A. Nagy, Z. Felföldi, and N. Johnson, "swap, swear and swindle incentive system for swarm," swarm-gateways.net/bzz:/theswarm.eth/ethersphere/orange-papers/1/sw%5E3.pdf, May 2016.
[15] L. Rizzo, "Effective Erasure Codes For Reliable Computer Communication Protocols," *ACM Computer Communication Review*, vol. 27, no. 2, pp. 24–36, April 1997.
[16] V. Buterin, "A note on data availability and erasure coding," github.com/ethereum/research/wiki/A-note-on-data-availability-and-erasure-coding, November 2017.
[17] C. Studholme and I. F. Blake, "Random matrices and codes for the erasure channel," *Algorithmica*, vol. 56, no. 4, pp. 605–620, Apr 2010.
[18] T. Le, J. S., and S. B., "Code sample: Intel isa-l erasure code and recovery," software.intel.com/en-us/articles/intel-isa-l-erasure-code-and-recovery, 2017.