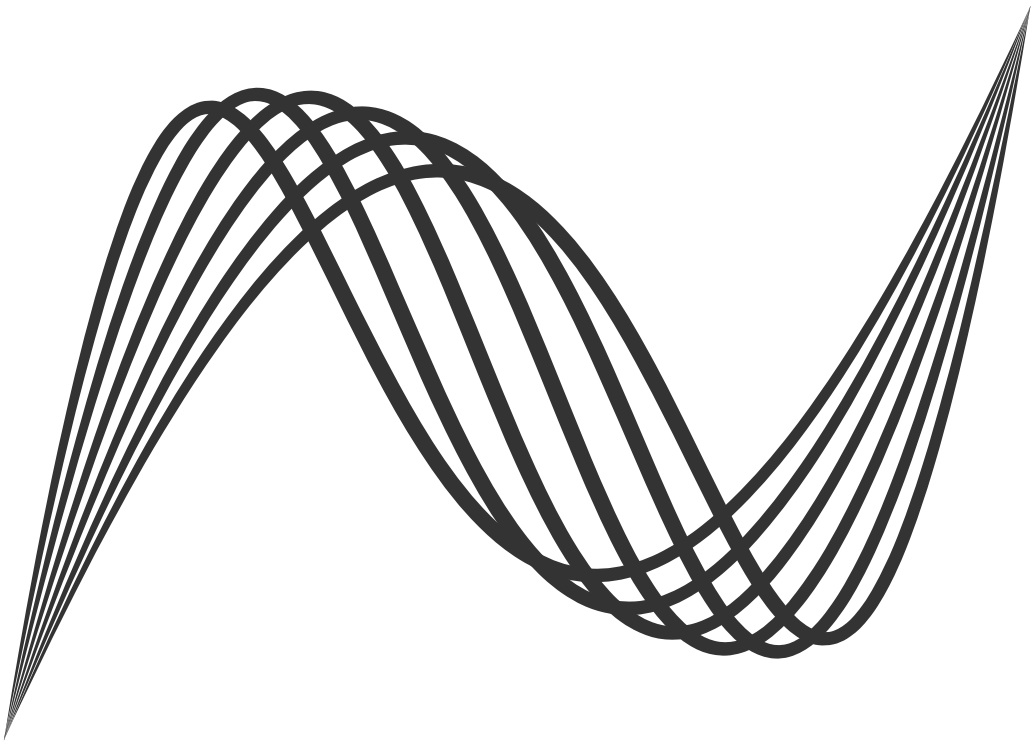

NEBULA NETWORK

Project Description

Version 0.8.7

If you're in a hurry there is a **TL;DR summary at the end**. Each chapter also contains a TL;DR subsection.

Table of Contents

Motivation

In its inception the internet was decentralized. Meanwhile the client-server model drove us slowly away from that premise. Today's clouds are centralized, they live on data centers usually controlled by a single entity. And more and more gigantic server farms are being built every day.

We believe the centralized model is just a hindrance along the way, that arose do to technical difficulties of the past. Decentralized systems are more robust, and can deliver better performance. Think of BitTorrent.

On the other hand we are seeing the rise of cryptocurrencies. They already embrace decentralization. Nevertheless they face some hurdles on the way to mainstream adoption. One such hurdle is the speed of the underlying technology – the blockchain. While the blockchain is impressive in many regards, the rate of transactions provided by the most popular cryptocurrencies based on it is abysmal. Currently Bitcoin authorizes less than 4 transactions per second and under optimal conditions the current protocol can scale only up to 7 transactions per second. Ethereum is able to handle only 15.

Another problem for many people is the perceived lack of intrinsic value. In fact not relying on a central authority by design, cryptocurrencies do not to have any major backer such as a central bank or a state. Most cryptocurrency balances can be perceived as just a number. Some even suggest that bitcoin is a pyramid scheme. There are attempts to mitigate this, for example OneGram and ZrCoin are commodity backed. However, physical commodity backing crosses the virtual and real world barrier. To cross this barrier an institution must warrant the possibility of exchange of the virtual tokens for the physical commodity. Such an institution needs to be trusted and thus it is hard to decentralize. The institution can be robbed of the commodities or can prove not to be trustworthy, and so it becomes a weak link that nullifies all the advantages and security of the decentralized system. The lager the market cap, the bigger such risks grow, as does the cost to alleviate them.

We believe that the above problems are interlocked and have a common solution. This document sketches out our solution.

TL;DR Motivation Summary

- The Internet is more and more centralized.
- Current cryptocurrencies are slow (low transaction rates).
- Current, truly decentralized cryptocurrencies are not backed by anything tangible.
- We want to solve those problems.

Our Vision

We want to democratize the cloud. While others build server farm cages for their clouds we strive to make a cosmic size cloud that spans the globe: a **NEBULA**.

The Stellar Spire Nebula as photographed by the Hubble Space Telescope in 2005. Public domain image. The soaring tower is about 9.5 light-years high.

To do that we want anybody to be able to join and contribute resources such as storage and computing power. There have been attempts to do such things in the past with varying degrees of success. For example the SETI project greatly enhanced its data processing capabilities using the SETI@home program which allowed volunteers to donate their CPU power. The SETI project is inspiring, however the computational network that it created is not universal. Many similar projects failed to attract nodes to their network as there was no clear incentive.

A while ago another network – Bitcoin – was created. Back when mining was easier all nodes had a strong incentive to be a part of the network which ensured its rapid growth. However the bitcoin network is also not an universal computational network. Mining is mostly calculating hash values to secure the network by a scheme called Proof of Work (PoW). For all other intents and purposed the computations are just a waste of energy. Furthermore as mining moved to ASICs, we observe progressing centralization and diminishing motivation for running full nodes on unspecialized hardware. Fortunately there is research into alternative blockchain securing schemes such as Proof of Stake (PoS).

Not only computation can be performed in a distributed network. Storage and data distribution are other key resources that can be distributed among the nodes of a network.

We propose an elegant solution where of binding such resources with a currency emission mechanism in a patent we're working on. A provisional patent application has already been successfully filed to United States Patent and Trademark Office (USPTO).

By binding currency emission with the amount of useful resources available within the network we effectively create the first cryptocurrency with intrinsic value without crossing the virtual to physical world barrier.

For it all to work smoothly vast amounts of microtransactions are needed. The biggest blockchains out there are orders of magnitude too slow for this. Basically the maximum speed a blockchain can operate on is determined by a very simple formula:

$$transaction\_rate = \frac{blocksize}{blocktime}$$

We want many miners (or stakers) to make the systems truly decentralized so we cannot push the parameters too high as the blocks cannot be propagated through the whole network fast enough which results in many orphaned blockchain forks. This is a fundamental blockchain property that some people try to mitigate in so-called private blockchains by limiting the number of consensus forming (minig) nodes and making sure the connections between them are extraordinarily fast. It is important to note that such an approach will be futile for public blockchains, if we insist that transactions from all over the world must be synchronized through a global blockchain, where every node verifies each transaction.

However for most applications not all transactions must be confirmed at the same speed. While most of us are ok with even an hourly delay when transferring funds to another continent, we wouldn't want to wait even minutes for a payment to be processed on the cash register. Having noticed that we came up with a novel Multichain Blockchain algorithm. Basically we have many

smaller blockchains that operate independently to a certain degree. To make them as secure as a single blockchian we interlock them by putting the hashes of the blocks into a masterchain (or a hierarchical system of masterchains if the network grows big enough). This interlocking is done with a delay. This allows us to desynchronize the speed at which different transactions are confirmed with various degrees of confidence, and greatly increases transaction throughput for most transactions (especially if the two parties have funds on the same sub-blockchain).

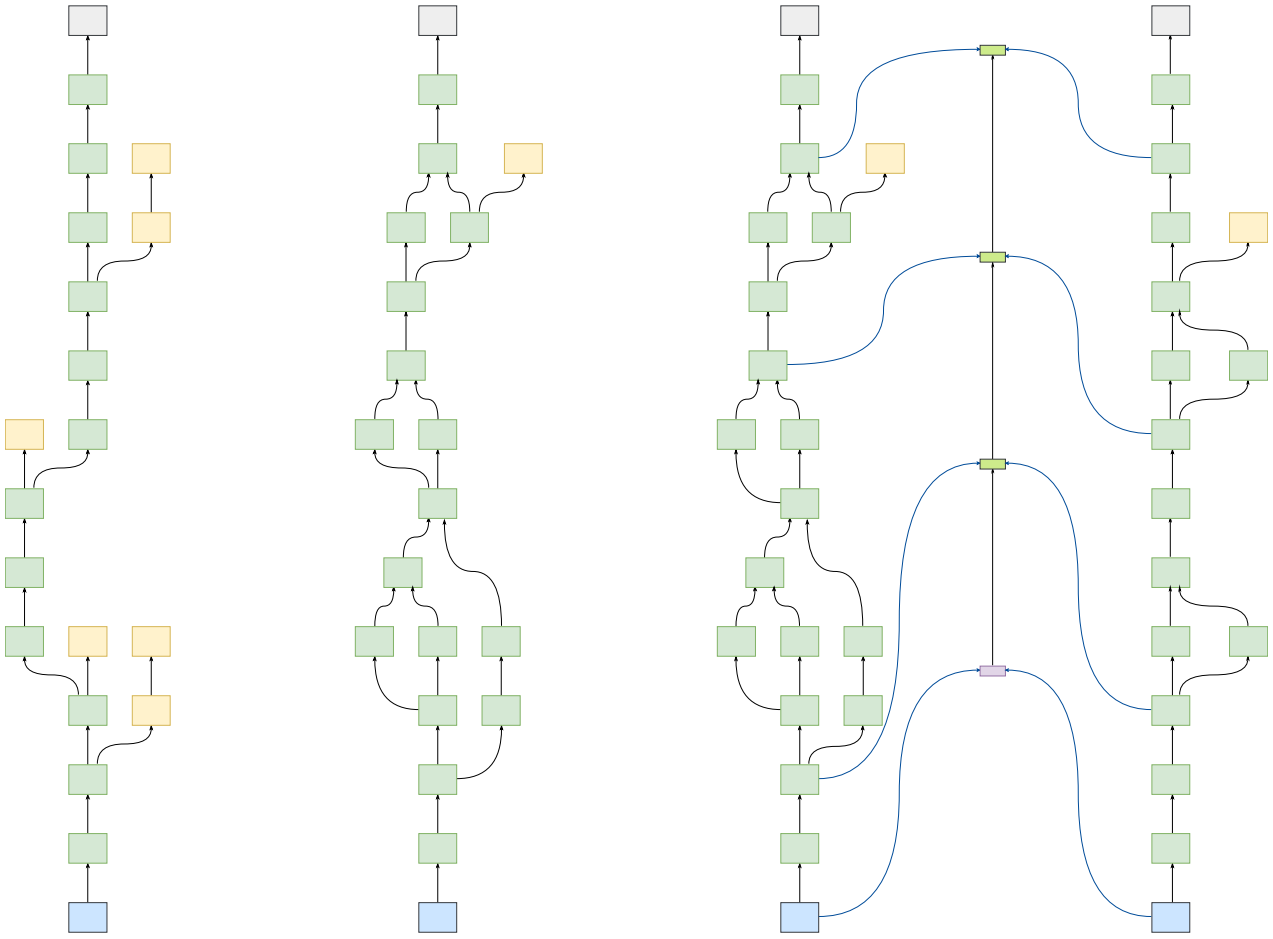

Comparison of a classical blockchain with a DAG based blockchain and a Multichain. Green blocks are the valid parts of the blockchain, yellow ones should be discarded, gray is the current block which should be followed, blue is the origin block. Multichain allows smaller chains that can go faster but are synchronized with a delay to provide similar security to a single blockchain.

To reach additional transaction-rate speeds we use an off-chain, second layer mechanism akin to the Lightning Network. Note that this can transfer funds between the sub-blockchains.

To summarize we greatly speed up the blockchain by

- a Multichain Blockchain approach detailed further in the Nebula Network whitepaper,
- an off-chain second layer mechanism based on the Lighting Network.

Note that the Lightning network as it is based mostly on off-chain transactions does not enhance the transaction rate of the blockchain itself.

Designing and testing a fast blockchain as a solid backbone for our cryptocurrency is a natural first step as it is very useful in itself. But a fast transaction rate is crucial to enable the sharing of resources with a low enough granularity to be useful.

To provide storage we divide the data into small chunks (shards), encrypt it and spread redundantly among other nodes of the network. Optionally it can also be signed to be able to control what the complying nodes which store the data do with it (for operations such as removal or sharing/access control). Prerequisites for a proof of retrievability should also be made to enable easier verification what pieces of the information are losing a comfortable amount of redundancy, so we can increase it (spread the data to more nodes) – which would help us the default redundancy much lower for the same amount of security. The data should not be labeled in a way that allows the uploading node to be identified which should make it impossible to impose not self governed censorship. Ideas from the TOR and I2P networks are being incorporated into the data spread algorithm.

To perform computations the user must compile their programs into LLVM bytecode. To execute it we use the PNaCl sandbox developed as open source by Google within the Chromium project. The chunks are then sent to other nodes of the network in a scheme described in more detail within our patent application. For security and robustness this should be also done with a certain degree of redundancy. Redundancy should also be done in such a way that the resource providing nodes do not know, and cannot control what other nodes were asked to provide spare (redundant) resources.

More detail is in our patent application and the Nebula Network whitepaper.

TL;DR Vision Summary

We want to achieve 3 things with the project:

1. Democratize the cloud. Both storage and computation.
2. Create a cryptocurrency with solid backing in resources such as computational power and distributed data storage.
3. Develop a much more efficient version of the blockchain with technologies based on the Multichain Blockchain, and Lightning Network.

Our Plan

Technologies Behind The Nebula

To be competitive our currency must with time provide sought after features of other major currencies. We also must use good tools for the job.

We strive to provide:

- A Multichain Blockchain, and a second layer protocol based on the Lightning Network.

- A sandbox based on PNaCL from the Chromium projects for the execution of LLVM code chunks to build the distributed computational network.
- Storage providing mechanism with anonymous data distribution.
- Smart contracts based on Solidity language introduced in Ethereum, with slight modifications for contract costs to avoid the miner's / verifier's dilemma related problems (i.e. not executing contract code for transactions to reduce the computational workload of mining new blocks, assuming that the others did so honestly already and found no invalid operations) and optimizations introduced by the RSK project.
- An optional global zero-knowlege-proof 'mixer' for simple transactions akin to the ZeroCoin or ZCash protocol to provide optional truly anonymous payments.
- A voting mechanism akin to Dash / Decred with a fund to sponsor further development in the direction determined by the community.
- A way to provide a digital identity which would be very useful for building social media on our platform.

We develop the project in C++. It is a language with many flaws but the newest versions (14, 17) gave it useful modern features, its performance is unmatched and the founders have great experience in that technology. It is also a proven technology in which Bitcoin is written, as well as PNaCl and many other pieces of technology which we will adopt into our project.

Roadmap

Secure intellectual property

To deter competition and prevent fragmentation we have a provisional patent application successfully filed to USTPO on a Cryptocurrency with Value Based on Electronic Resources Made Available to Peer Nodes and there is a second one being prepared on our advancements of the blockchain technology.

Build up the team

We strive to build and global international team of the best people in the business. On our tech team watchlist are people who originally developed PNaCl at Google, authors of LLVM, authors of various BitTorrent clients, contributors to the TOR and I2P projects, as well as graduates from top Polish IT Universities and cryptocurrency veterans.

Even more urgently we are expanding our marketing and business team as they are obviously critical to the success of the project. We see several opportunities on that front as well as we talk with various parties.

Market the project and collect funds through an ICO or otherwise

An ICO is currently a great way to get funding for project such as ours.

Deliver a scalable and fast blockchain technology

This should be the first step as it is useful in itself, and is the foundation for the rest of the project.

Deliver distributed data storage

Data storage solution shall be developed in parallel to the distributed computational engine, but we expect the storage to be simpler to secure, and it will not require us to provide such an extensive API and documentation as computation, so we plan to roll storage out first.

Deliver distributed computing

We will harness the PNaCl sandbox and send chunks of LLVM bytecode to computing nodes. We also need to provide an API for application developers and documentation including tutorials.

We think deep-learning (think a 'TensorFlow backend') and distributed rendering (integrated into an open source program such as Blender/Cycles) could be great showcases for our technology.

Deliver easy web-site deployment

We think that once both storage and computation are deployed most people will want to host Websites on our platform. We strive to make porting websites that are based on Node.js & MongoDB and PHP & SQL trivial to port and deploy onto our platform.

Construct distributed social media

We think the world needs an uncensored distributed social media platform. Imagine Bit Torrent and Facebook combined. That would be the ultimate use for our platform.

TL;DR Plan Summary

- Secure IP.
- Build up the team. The project inspires people so we hope to be able to attract top talent for all around the globe.
- Go through a croundfunding process – an ICO.
- Gradually rollout key features. In order: fast blockchain, distributed storage, distributed computation, easy website deployment, ...
- Gradually copy the features of other popular cryptocurrencies as well – Solidity smart contracts, community controlled development fund, optional transaction anonymization, digital identity
- Use proven technologies that are able to deliver top performance; C++, PNaCl and parts of other open source projects.

Competition

Filecoin & IPFS

Those projects also realize the potential of distributed storage with blockchian based payments. The people behind it are affiliated with archive.org (Wayback machine) and believe the storage should be permanent (you cannot delete once you upload). We do not believe that corresponds well with businesses needs or is compatible with the “right to be forgotten” introduced some time in the European Union.

Golem

Golem is being developed by a cool team in Warsaw Poland. Golem also realizes the need for democratizing computation. In that respect their project is similar to ours. They had a very successful ICO some time ago (late 2016) collecting over 8 million dollars in under half an hour.

Golem lives on the ethereum blockchain. As we mentioned earlier the ethereum blockchain is rather slow in itself, and thus a big part of the Golem project was figuring out how to do quasi microtransactions that could be used for paying for the computations. We do not think that it can be done very efectivley and that it is worth the effort, as the Ethereum blockchain was not designed with projects such as Golem in mind (although the Ethereum organization is trying to help Golem).

Golem does not really offer a solution for storage, but computation is often performed on a set of data so we believe a solution for both is needed to succeed. To solve that problem Golem wants to integrate with IPFS – but exactly how this will be achieved is not clear, neither is how Golem wants to deal with the fact that IPFS is permanent (which does not lend itself well for temporary data). Furthermore data storage can allow for better growth in the beginning as distributed computation starts to show its potential only when many nodes are already the network, so that together they clearly outperform other solutions. To our knowledge the project falls short in that regard.

There also seems to be no fraud prevention that would protect the computation requesting nodes from bad results except for white and blacklisting of applications.

According to the Golem whitepaper there will be *“no token creation, minting or mining after the crowdfunding period”*, thus the supply of the tokens is unlinked form the supply of the computing power present in the network. Many believe that to be a completely flawed economic model.

Finally it seems that Golem applications are bound to the x86 architecture as the used Docker sandbox does not provide architecture independence. In the future Golem apps could be cross compiled but that would introduce many additional problems. By using LLVM code Nebula Networks solves this problem and makes architecture specific vulnerabilities harder to exploit.

Storj

Storj is a platform for decentralized storage. It is based on the Kademilia protocol which gives it a solid foundation. It is written in Javascript and Node.js. Unfortunately Storj has a few problems:

- Billing goes through their centralized servers and is not cypto-currency based.
- Due to implementation problems they deemed usage of proofs of retrievability to inefficient.

Siacoin

Siacoin embraces distributed storage, and payments via cryptocurrencies, however it does not tie the supply (emission) of the coins with the supply of storage that is provided; Sia whitepaper states *"The supply of siacoins will increase permanently, and all fresh supply will be given to miners as a block subisdy"*. Sia is implemented in Go. Sia is a product of Nebulous Incorporated. Sia is intended to become a primary source of income for the company. Sia offers only storage.

Traditional clouds

The cost of operation will be cheaper for those due to the effects of scale, nevertheless they have to factor in the cost of the machines. For many people joining our platform the device they provide would be purchased anyway, so they can think of the hardware as being free.

Also traditional clouds cannot provide the level of trust that our platform can, being controlled by a single entity, that often also wants to mine user data (for example to run targeted ads or to comply with NSA requests)

In a way traditional cloud solutions compare to Nebula Network like a public transport system compares to Uber.

Currencies backed by commodities

As mentioned earlier currencies backed by physical goods need an institution that nullifies most cool things about a cryptocurrency. Examples include:

ZrCoin

ZrCoin is just a way to invest into a Zirconium Dioxide recycling process "invented by Russian scientists". They successfully completed an ICO in 2017.

OneGram

One gram is a cryptocurrency targeted at the middle eastern markets and strives to be a currency compatible with the Islamic Sharia law. As of June 2017 it has an active log running ICO. According to the OneGram whitepaper *"While Bitcoin and many other cryptocurrencies are much closer to the Islamic definition of money than modern fiat money, they still fall short, as they are not backed by any tangible real world asset."* Each transaction of OneGram Coin (OGC) generates a transaction fee of 1% up to a maximum of 1 OGC, which is reinvested in more gold (net of admin costs), thus increasing the amount of gold that backs each OneGram – assuming the stakers reward is smaller than the amount reinvested in gold – which OneGram does not specify. There is no source code, and technical information is scare.

TL;DR Competition Summary

Similar projects are generating a lot of enthusiasm, even though they are riddled with problems, and none provide a complete solution consisting of a new fast blockchain, distributed storage, distributed computation, and tying the currency emission to available resources – each of the projects we know of innovates only in one of the above areas. We can be the first comprehensive and thus truly working solution.

TL;DR Overall Summary

- We are developing a democratic cloud where anybody can join to offer resources such as storage and computational power. We use cryptography and redundancy to make it secure for all parties.
- We incentivize people to provide such resources by giving them cryptocurrency.
- The main emission mechanism of the cryptocurrency is tied to the amount of provided resources, thus backing the value of the currency with a sought-after asset, but without crossing the virtual-real world barrier – in a fully decentralized way that does not require trust to any central institution.
- We use advances in blockchain related research to make the cryptocurrency transaction rate fast enough to support vast amounts of microtransactions which are crucial for providing the resources at a low enough granularity.
- We will further build infrastructure on top to provide distributed uncensored and reliable social-media with embedded file sharing and other services further along the way.
- We already filed some provisional applications to USPTO related to the technologies within this project, more coming.

Think big. Think **Nebula Network**.

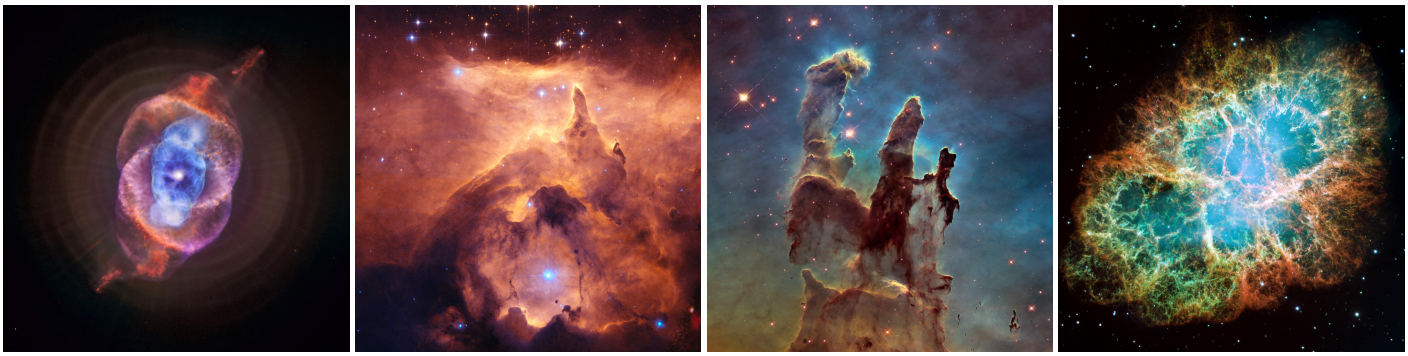

From left to right: Cat's Eye Nebula, Lobster Nebula, Pillars of Creation fragment of the Eagle Nebula, Crab Nebula. Public domain images.