

# A Brief History of P2P Content Distribution, in 10 Major Steps

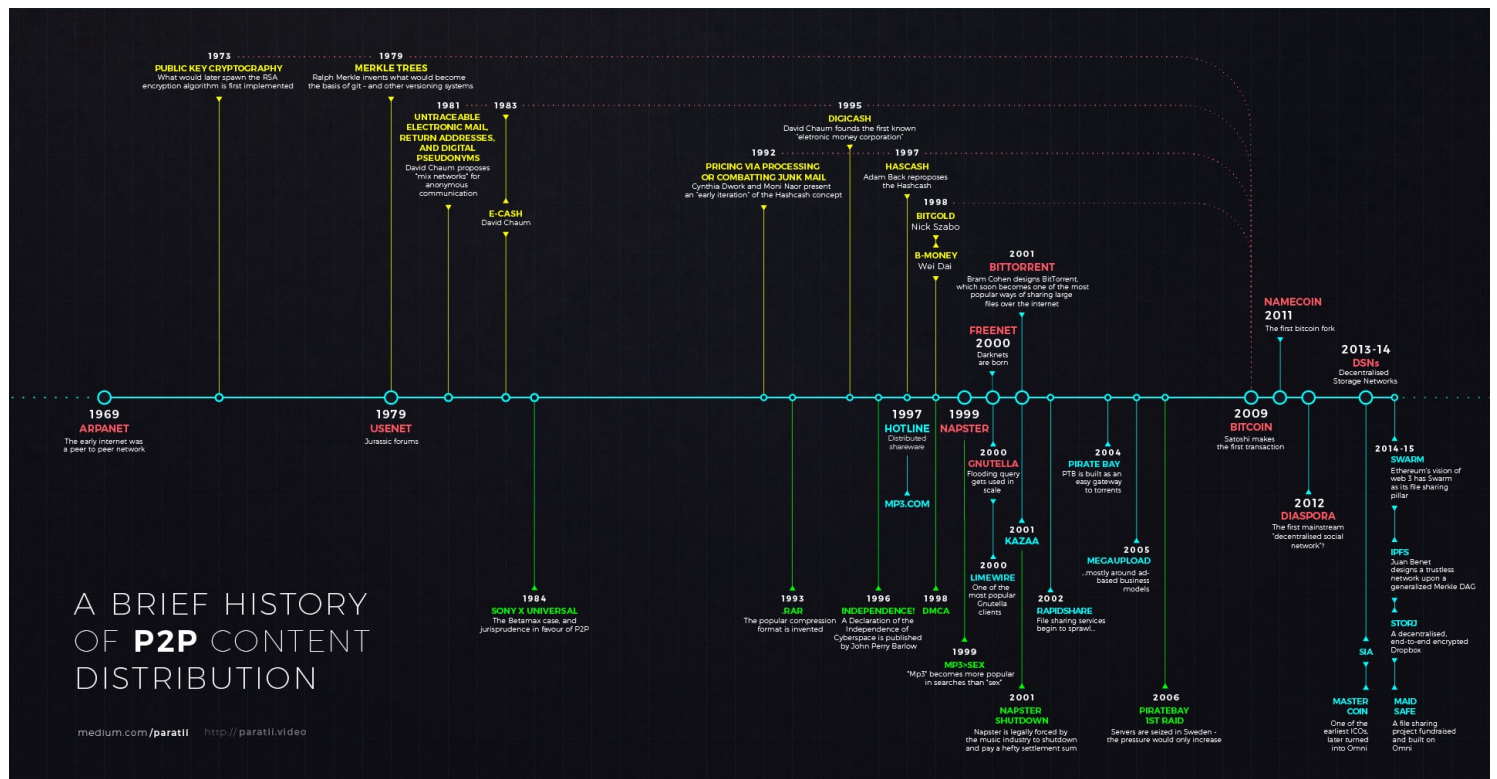
Peer to peer networks build upon decades of cryptographic research and battleground experimentation. It's tough for any particular initiative to be labeled as uniquely "pioneering".



Paratii Follow

Oct 25, 2017 · 9 min read

*This is the first of a 3-piece series on decentralised video distribution. The second piece is coming out this week. Subscribe [here](#) to receive it in your inbox.*



Please tell us of any corrections you find should be made, [via telegram](#)

The paradigm shift from a client-server architecture to peer to peer has been

explored from many different perspectives. A useful metaphor is that of the change from slave-based Rome to feudal societies in the West. No longer able to grow in space, through conquests, and facing linearly increasing costs *per-resource unit*, the empire-machine becomes inefficient, and vulnerable to competition.



Thomas Cole, on the Roman Empire. They also had trouble with scaling.

Alternative social systems who better put to use the potential of their internal resources thus begin to sprawl. Why use slaves for conquering only, when they can seed and grow crop, police the land, and much more? Extensive development gives room to intensive development, and feuds put their *resource-units* to cooperate and produce together, while everyone that joins in not only consumes scarce resources, but also provides them.

The narrative is roughly analogous to the rationale behind peer to peer file sharing systems. In a server-based architecture, the more demand for a file, the more bandwidth cost it consumers. In a peer to peer system, the more a file is requested, the more nodes are seeding it, lowering the delivery cost per-file distributed. The role of audience in the equation is inverted.

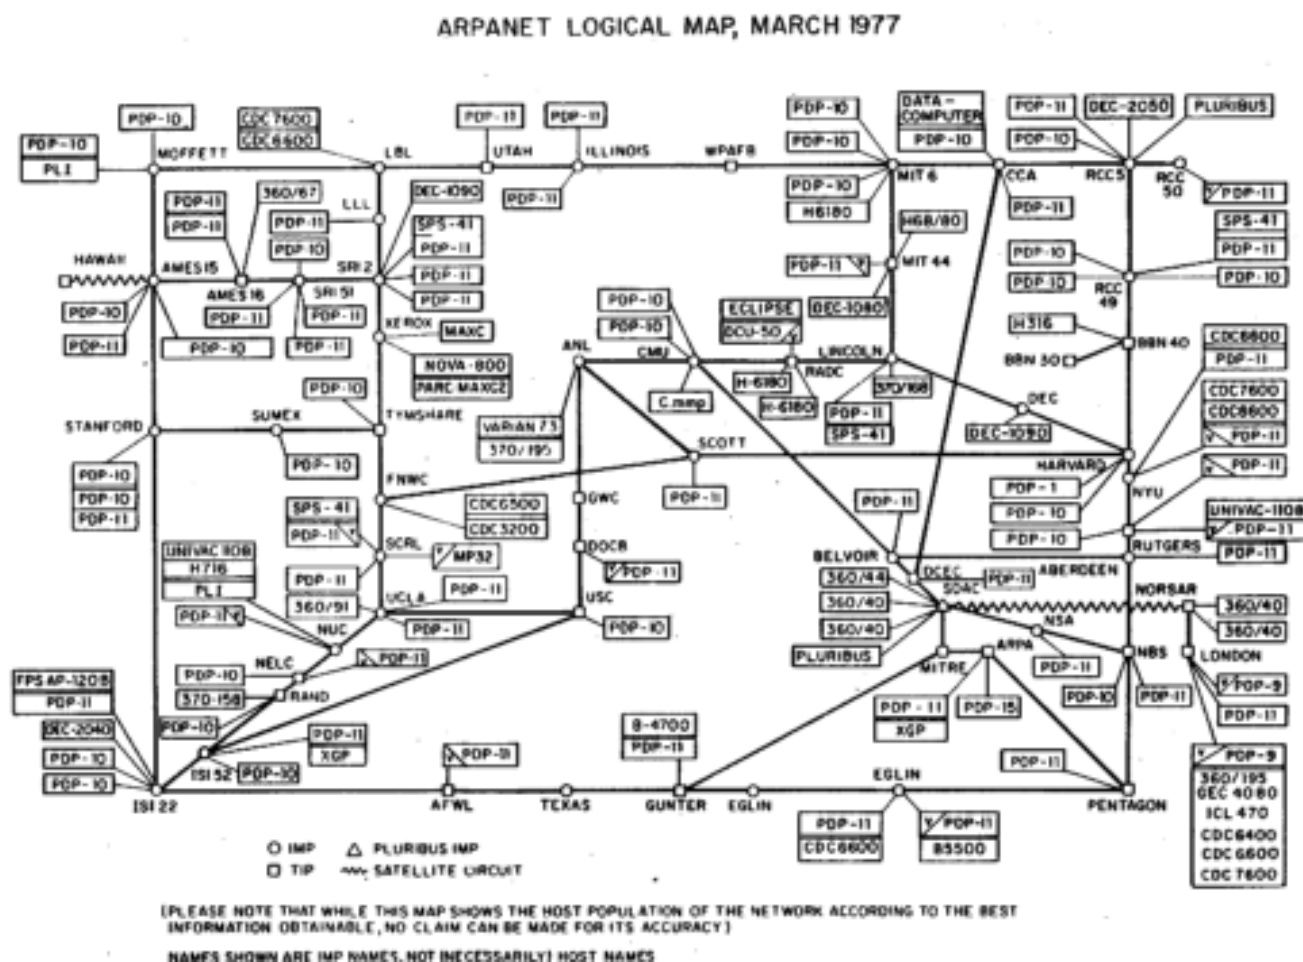
The peer to peer movement is rooted in the early internet, and has given birth to countless protocols and applications that, on the most extreme cases, redefined our way

of consuming entertainment. P2P can be self-scalable, censorship-resistant, anonymous — and the robustness of present implementations is the product of incremental evolution.

Some of the recent work around **decentralised storage networks** (as the Protocol Labs team labels) is extremely novel, and blockchain-enabled incentivisation indeed resparks the dream of fully unstoppable file distribution. Nevertheless, it's useful to remember that all this is built upon some inestimable heritage that goes back half a century in time.

## A non-exhaustive selection of notable initiatives

The history of p2p is permeated with hundreds of initiatives that could make for a comprehensive list - this is only meant to be a gentle introduction to a field of research that has been increasingly spawning global socioeconomic experiments. The past tense is used for narrative purposes - most of the networks mentioned below are still up and running. Even more interesting: most of the *people* who built them are around us.



This definitely looks distributed.




## 1 - 1969 - The ARPANET

The ARPANET originally connected the UCLA, Stanford Research Institute, UC Santa Barbara and the University of Utah not in a client/server format but treating them as equal computing peers.

Some of the early popular applications of the internet like FTP and Telnet followed themselves client/server architectures, but since hosts could act as servers to other hosts, symmetric usage patterns emerged.

## 2 - 1979 - Usenet

Usenet was developed by American graduate students, based on the Unix-to-Unix-copy protocol (UUCP). Through it, a Unix machine could automatically dial another, exchange a file, and disconnect, in what resembled a bulletin board system (BBS) that's somehow precursor to the forums and feeds we have today. Usenet arguably gave birth to terms such as "FAQ" and "spam" :)

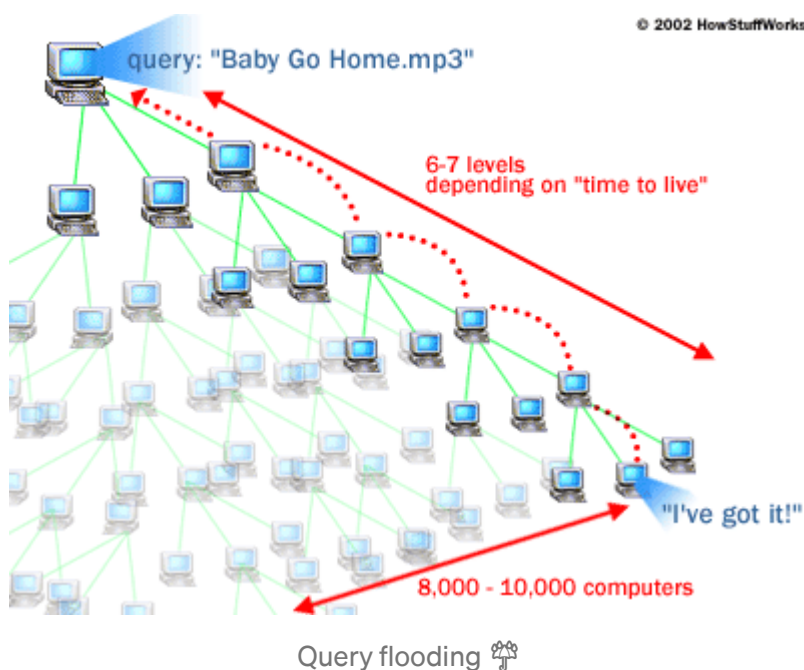


 Napster had a central index that told peers who to talk to

### 3 – 1999 - Napster

Throughout the 80s and 90s, the client-server model flourished since the power of CPUs available to consumers was still small (albeit rising). Most file transfers happened over landline telephone, with FTP and USENET gaining usage along the period, and the IRC being invented in 88. On the late 90s, new data compression technologies (MP3, MPEG) came into mainstream use and were probably the last straw.

This Internet zeitgeist began shifting back to peer to peer with the introduction of Napster, developed by Shawn Fanning while still a freshman at Northeastern University. Users downloaded a free program that searched the local disks of neighbouring computers for MP3 files and then were able to download these directly from one another. In less than one year, Napster had over a million members. In less than two, Metallica filled a lawsuit against the company. In less than three (July 2001) it was shut down after legal suffocation and a failed attempt to become a pay-based service.



### 4 - 2000 - Gnutella

Napster still relied on the operation of central indexing servers, and thus was susceptible to shutdown. A new breed of file sharing spearheaded by Gnutella eliminated such vulnerability by allowing users to find each other and connect remotely. By employing a query flooding model, the protocol made each search be broadcast successively to other machines in the network (which, in the case, was significantly less efficient than

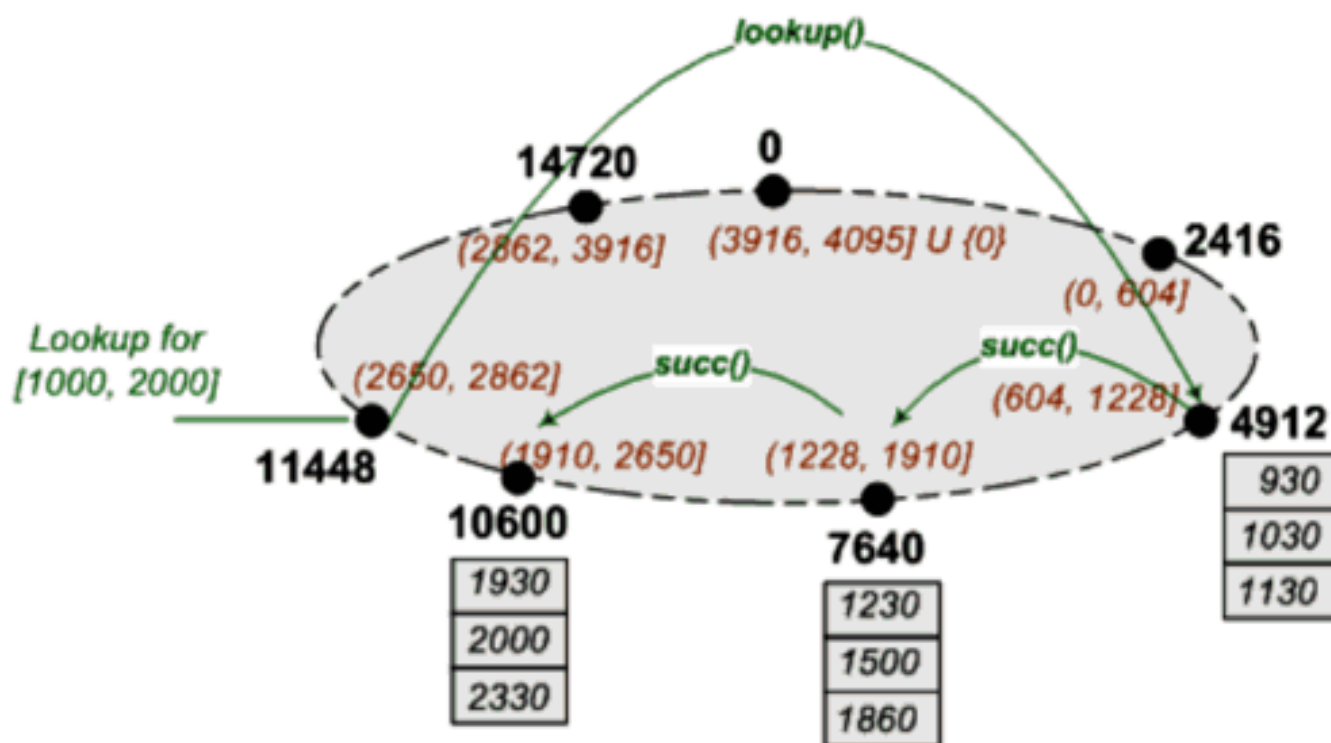
querying the central indexer). Another difference between Gnutella and its predecessors was the amount of clients now available to run the protocol. LimeWire, for example, was one among many Gnutella clients.

## 5 – 2000 - Freenet

Freenet brought a major improvement in regards to user anonymity, inaugurating what would later be labeled the “darknet” category. It made users store encrypted snippets of files, connecting them only through intermediate computers that passed forth and back requests without knowing the contents being sought. The design resembles how routers on the Internet exchange packets without knowing anything about them.

## 6 - 2001 - Bittorrent

The protocol created by Bram Cohen allowed peers to communicate directly over a TCP port, but relied on central *trackers* to record the location/availability of files and to coordinate users. Vuze (ex-Azureus) was the first BitTorrent client to migrate to a trackerless system by implementing a distributed hash table (DHT) for peer discovery - other soon followed.

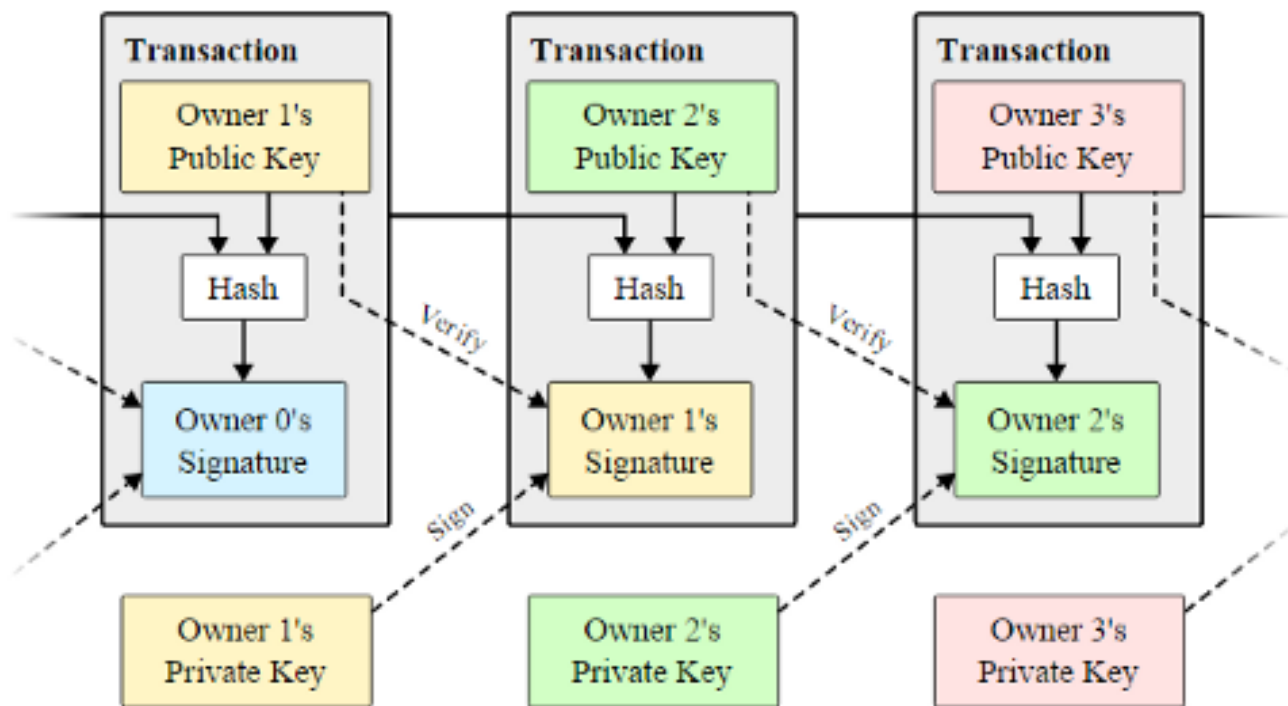


The illustration of a lookup in a DHT

From then on, BitTorrent emerged as a “two-sided” creature: protocol-following clients came with a DHT node that maintained a routing table with the contact information of some other nodes. When there was a search for a torrent among peers, the node used a distance metric to compare the infohash of the torrent it was seeking with the IDs of the nodes in its own routing table, iteratively finding ones that were closer and closer to the target infohash. Without the need for a centralised indexer, the network could persist with nodes updating their own routing tables indefinitely, even if legal authorities acted against specific entities or servers.

Such design allowed for new, interdependent networks to be created for each torrent, moulding subgroups according to the popularity and reach of their files. Even with Napster being shut down in 2001, peer to peer networks’ adoption kept growing at full speed. In August of the same year, more files were downloaded over Gnutella, Audiogalaxy and iMesh than what was registered by Napster on its most active month ever.

## 7 - 2009 - Bitcoin



A closeup on bitcoin transactions.

Bitcoin was not designed with file sharing in mind, but eventually spawned an entire new class of p2p storage frameworks. The blockchain is a distributed registry with a different purpose than that of DHTs: Satoshi wanted to store with each node an ever-growing transaction record without any chance of tampering and revision; DHTs aim at providing an efficient (in terms of lookup time and storage footprint) structure to divide data over a network, where immutability isn't the main priority.

What miners probably didn't expect was that their core activity would soon be abstracted into the generic category of *consensus mechanisms*, and applied to use cases far different from that of storing and transacting financial value.

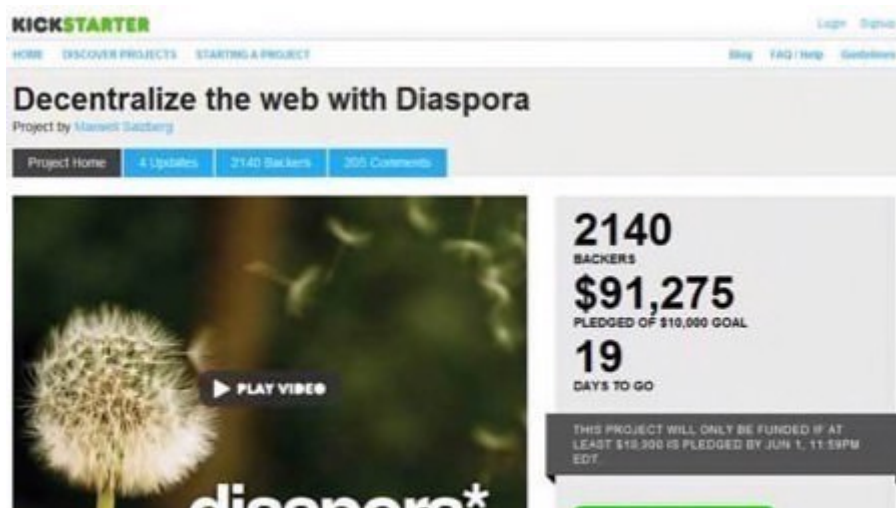
## 8 - 2011 - Namecoin

Namecoin was born out of the will to register data on the blockchain and make it less application-specific - it's flagship usecase is the top-level domain .bit. The first bitcoin fork, it's roughly a key/value pair registration and transfer system based on its fathering technology. It might be fair to say Namecoin represented the first non-monetary approach to a blockchain.

Vitalik talks enthusiastically about it in an early text on Ethereum for the Bitcoin Magazine, where he also mentions the potential of DNSs (see below).

## 9 - 2012 - Diaspora

Diaspora marketed itself as an open source personal web server and decentralised social network. It was funded with over U\$ 200.000 via Kickstarter in 2010, and released a short-lived consumer alpha shortly after, but reached a stable community release only in 2012.







If only ICOs arrived a couple of years earlier...

On it, users set up their own “pod” to host content, and interacted via a desktop client. The concept always revolved around aggregating content from other social networks, allowing posts to be easily imported from Facebook, Twitter, and so on but many of the required APIs were never made. The project whirl winded for some years before being embraced by Eben Moglen and the Freedom Box Foundation.

Diaspora sparked on the media an anti-Facebook reaction that’s similar to what we’ve been seeing with regards to fresher, recently tokenised social networks.

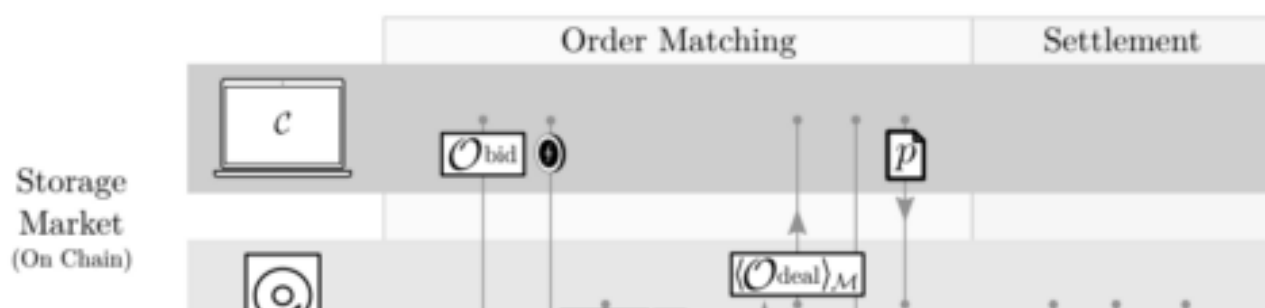
## 10— DSNs (Filecoin/Swarm/Sia/Storj/Maidsafe)


The idea behind decentralised storage networks is basically to turn cloud storage into an algorithmic market, by incentivising “miner” nodes to share storage space in order to be rewarded in a native token. **Incentivising** is the key word here. In BitTorrent, we already had some kind of content-addressed system where nodes seeded the files of others. Also, there was already a basic tit for tat mechanism for punishing badly behaving nodes.

There’s arguably technical improvements in, say, BitSwap (IPFS’ data trading module) over BitTorrent: one is not being limited to the bundle of data in the original torrent, but is rather able to fetch blocks from completely unrelated files, and to precisely identify what is being provided / requested, regardless of whether it’s a piece of a file, an input from a dataset, or a huge set of files. The big improvement is in how this is achieved (2):

(1) Providing the adequate data structures and network to interact with.

**(2) Creating a decentralised market fuelled by a native token for the storage and retrieval of chunks/files.**





An overview of the Filecoin (IPFS) scheme.

(1) has a lot of live examples. (2) has few working implementations, and many interesting projects now reaping the fruits of cutting edge research supported by heavy funding.


It's important to note that fully decentralised solutions, like the ones being developed to run on IPFS and Swarm, for example, are not yet close to being functionally implemented (although the underlying networks are there).







*Proof-of-retrievability, proof-of-storage, proof-of-spacetime* are hot research topics that basically trace back to the same issue: how to reward nodes for disk space and bandwidth natively, without much overhead regarding the proofs they should present, and in a trustless, effective manner?

Building upon such protocols, applications can lower distribution-related costs and are induced to allow for direct negotiation between peers, which leads to decentralised value flows and healthy feedback loops. The success of DSNs is pivotal for the dissemination of dApps and for accelerating the gradual transformation of business models into network incentive models.

If we were to classify peer to peer file sharing in “epochs”, 4 could be a good number. 1) we had Napster, which still relied on a central server to coordinate lookups within the network. 2) query flooding came into place with Gnutella, overloading nodes in a given network but eliminating central indexers. 3) DHTs gained adoption with the BitTorrent protocol, distributing the indexes themselves. 4) DSNs bring the promise of off chain token-based markets for storage and distribution.

. . .

This is an objective recap in order for us to see more clearly the grounds for what comes afterwards. The timeline presented here goes until 2015 only. Between then and now, hundreds of projects have started building tokenised applications and tools upon DSNs. On the next article, we'll have a look at some of these, with a particular focus on what touches the video market. Meanwhile, 

-  Follow @Paratii on [Twitter](#)
-  [Subscribe to our newsletter](#)
-  [Like us on facebook](#)
-  [Chat on Gitter](#)
-  [Track development progress on Github](#)
-  [Read more on Medium](#)

Powered by [Upscribe](#)

[Blockchain](#)[Peer To Peer](#)[Computer Science](#)[Technology](#)[Infographics](#)[About](#) [Help](#) [Legal](#)



Get the Medium app

