

# Helia & Performance

How to wring the most out of your deployment





## IP Stewards Team

Maintainer

- Helia
- js-libp2p
- Many supporting libraries

**Alex Potsides**

Protocol Labs



**JavaScript is slow?**





**Yes, if used badly**

- **CPU**
- **async**

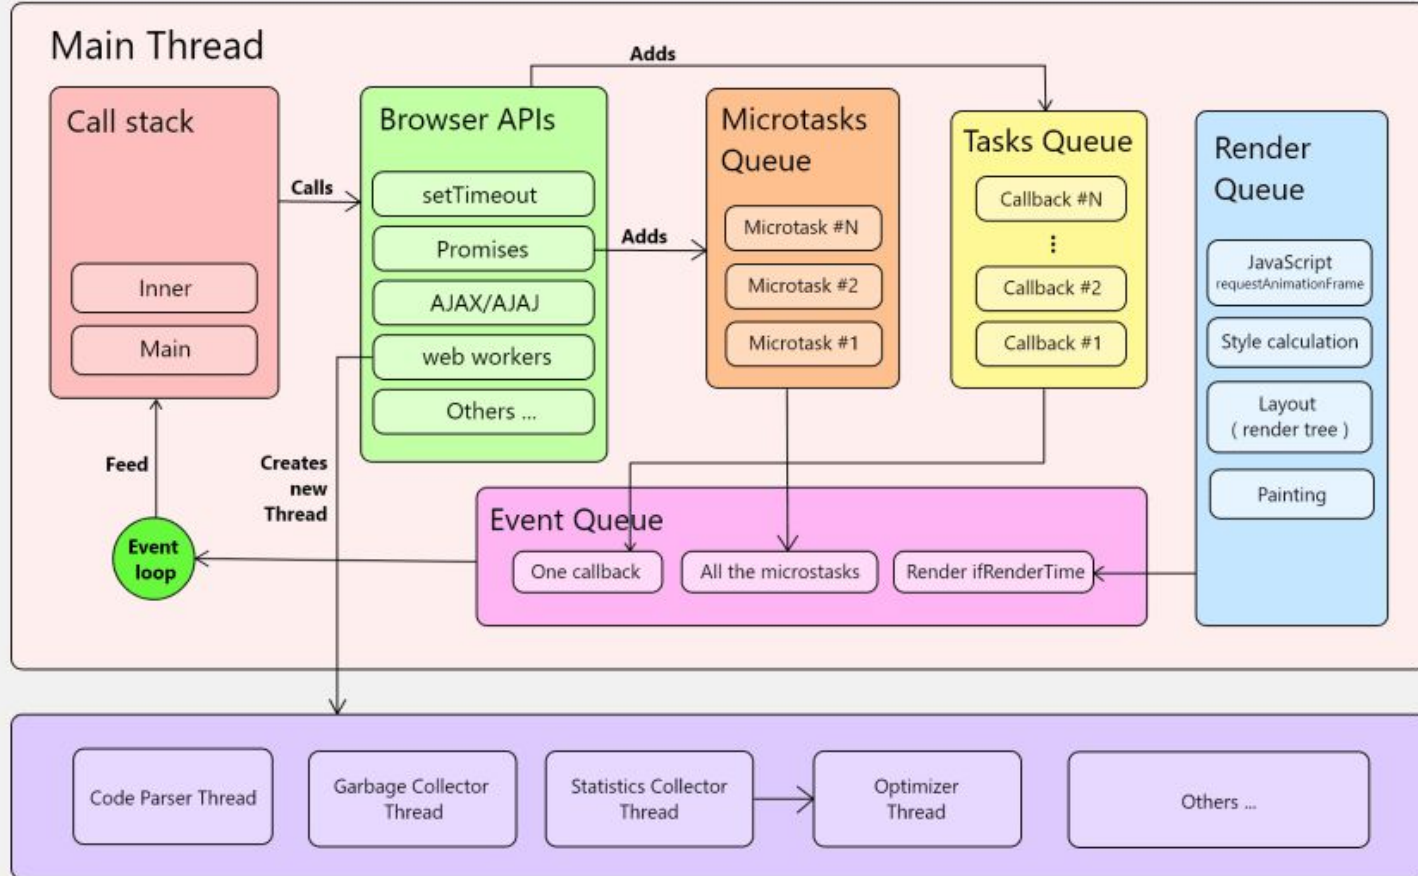


**Async is not free**





# Web Browser ( chrome / firefox )



**Demo**



# Storage matters







- **I/O is a massive bottleneck**
- **Blockstores can be slow**





## S3 Block/Datastores

# Best practices design patterns: optimizing Amazon S3 performance

[PDF](#) | [RSS](#)

Your applications can easily achieve thousands of transactions per second in request performance when uploading and retrieving storage from Amazon S3. Amazon S3 automatically scales to high request rates. For example, your application can achieve at least 3,500 PUT/COPY/POST/DELETE or 5,500 GET/HEAD requests per second per partitioned prefix. There are no limits to the number of prefixes in a bucket. You can increase your read or write performance by using parallelization. For example, if you create 10 prefixes in an Amazon S3 bucket to parallelize reads, you could scale your read performance to 55,000 read requests per second. Similarly, you can scale write operations by writing to multiple prefixes. For more information about creating and using prefixes, see [Organizing objects using prefixes](#).

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizing-performance.html>



## S3 Block/Datastores

# Best practices design patterns: optimizing Amazon S3 performance

[PDF](#) | [RSS](#)

Your applications can easily achieve thousands of transactions per second in request performance when uploading and retrieving storage from Amazon S3. Amazon S3 automatically scales to high request rates. For example, your application can achieve at least 3,500 PUT/COPY/POST/DELETE or 5,500 GET/HEAD requests per second per partitioned [prefix](#). There are no limits to the number of prefixes in a bucket. You can increase your read or write performance by using parallelization. For example, if you create 10 prefixes in an Amazon S3 bucket to parallelize reads, you could scale your read performance to 55,000 read requests per second. Similarly, you can scale write operations by writing to multiple prefixes. For more information about creating and using prefixes, see [Organizing objects using prefixes](#).

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizing-performance.html>



## S3 Blockstores

Default sharding strategy:

**CID(bafyfoo...abcdef)**

**-> /fe/bafyfoo...abcdef**

<https://github.com/ipfs/js-stores/blob/main/packages/blockstore-s3/src/sharding.ts#L20>



## Other stores are available:

- **MountDatastore**

- Combine multiple datastores separated by key prefixes
- You might use a memory datastore for the /peers prefix - all peer data will be stored in memory, everything else on disk

<https://npmjs.com/package/datastore-core>

<https://npmjs.com/package/blockstore-core>

## Other stores are available:

- **TieredDatastore**

- Combine multiple datastores
- Write to all, read from whichever is fastest?

<https://npmjs.com/package/datastore-core>

<https://npmjs.com/package/blockstore-core>





**Maybe invent your own?**

- **HotContentBlockstore?**

- Cache the 1000 most requested blocks in memory?

**It's up to you!**

**<https://npmjs.com/package/interface-datastore>**

**<https://npmjs.com/package/interface-blockstore>**



# Content routing matters





- **Resolving content takes a long time**
- **Configure delegates**
  - **Streaming (fast)**
    - @libp2p/ipni-content-routing
  - **Non-streaming (slow)**
    - @libp2p/reframe-content-routing
    - @libp2p/delegated-content-routing

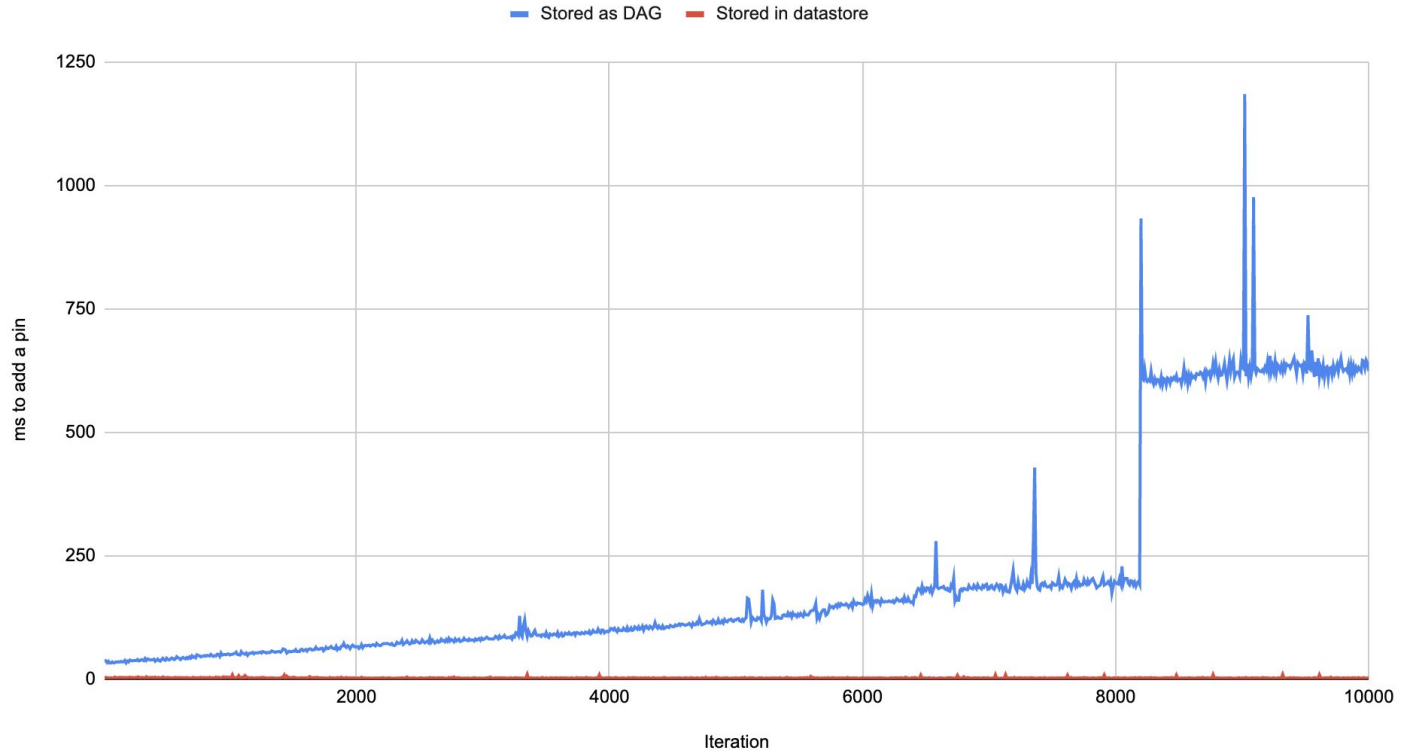


# Pinning & garbage collection



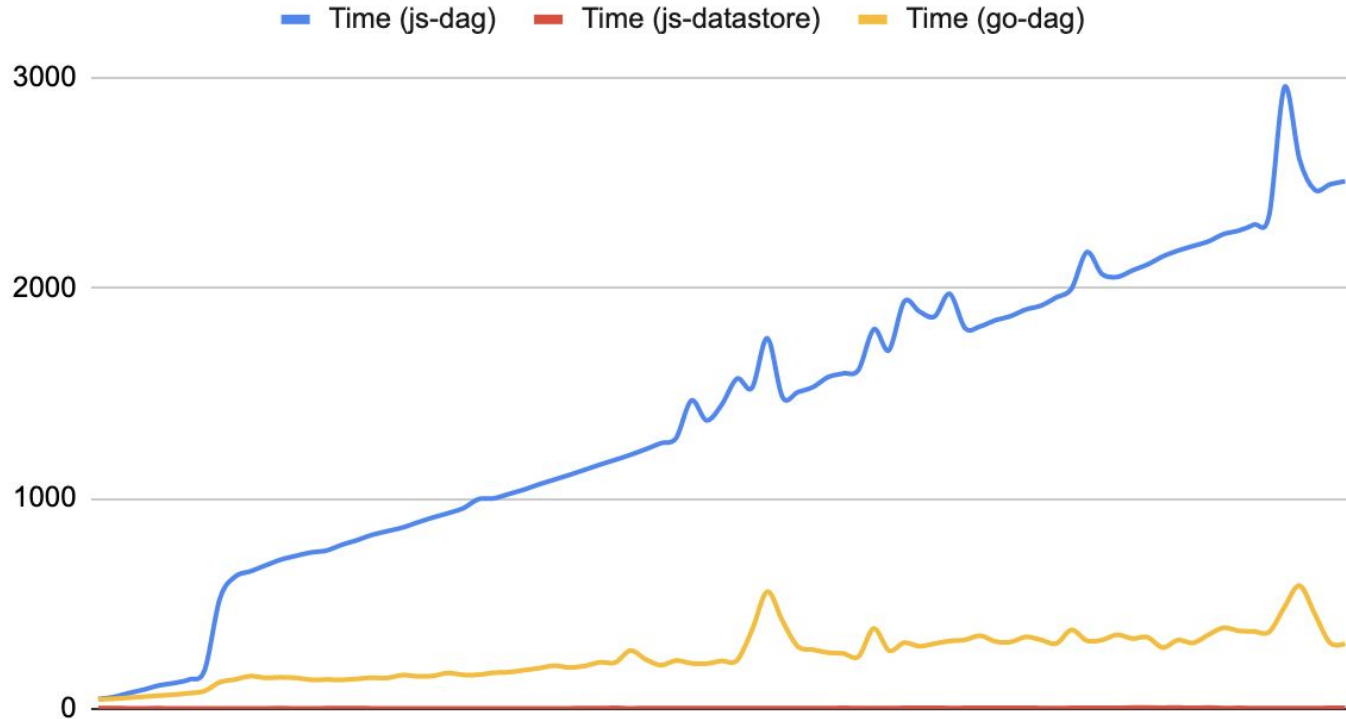


## Stored as DAG vs stored in datastore






## Time to add a pin





## 100k pins stored in datastore

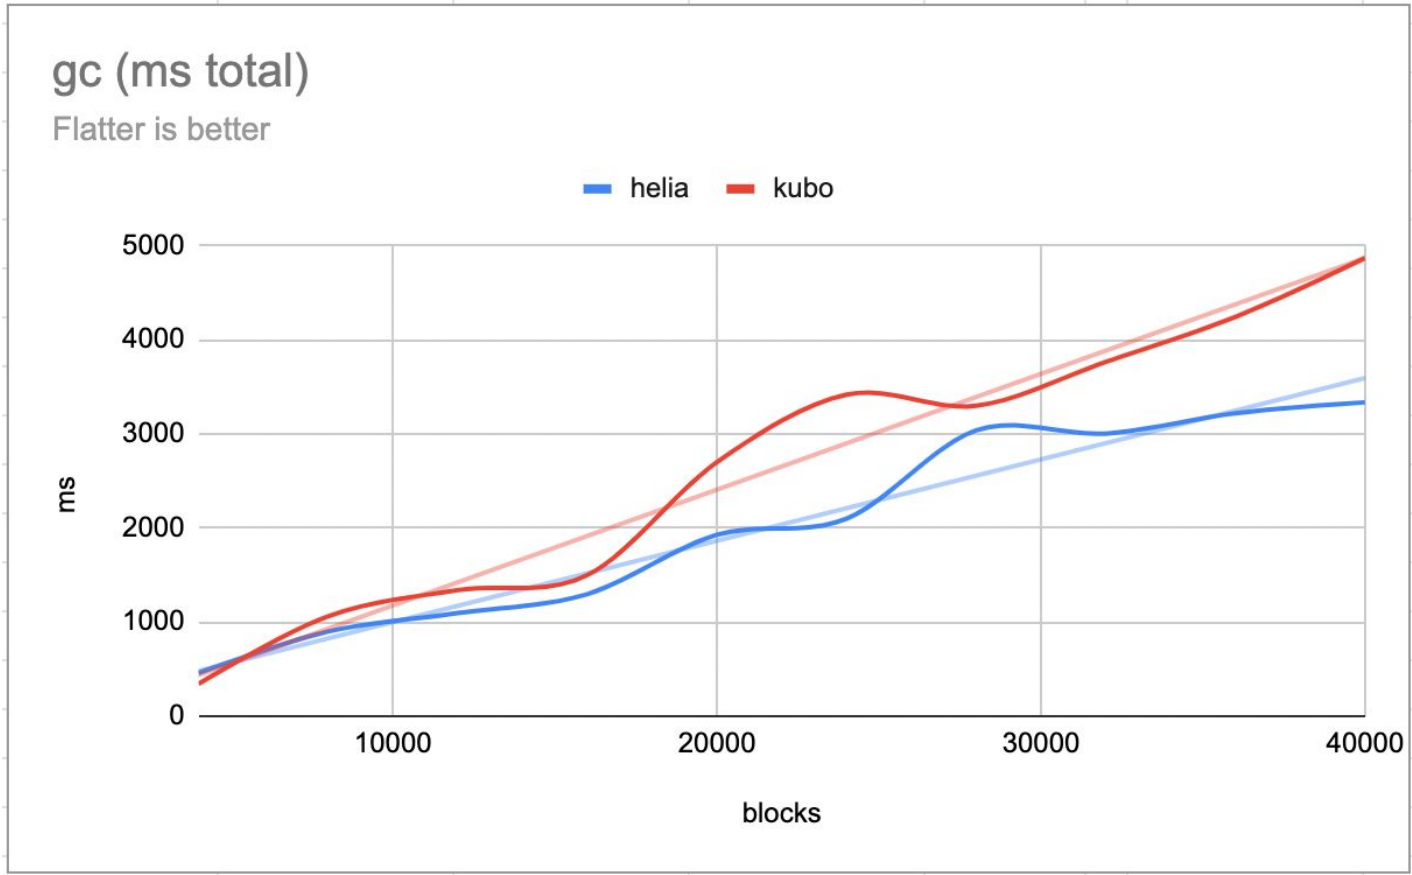


**Until now**

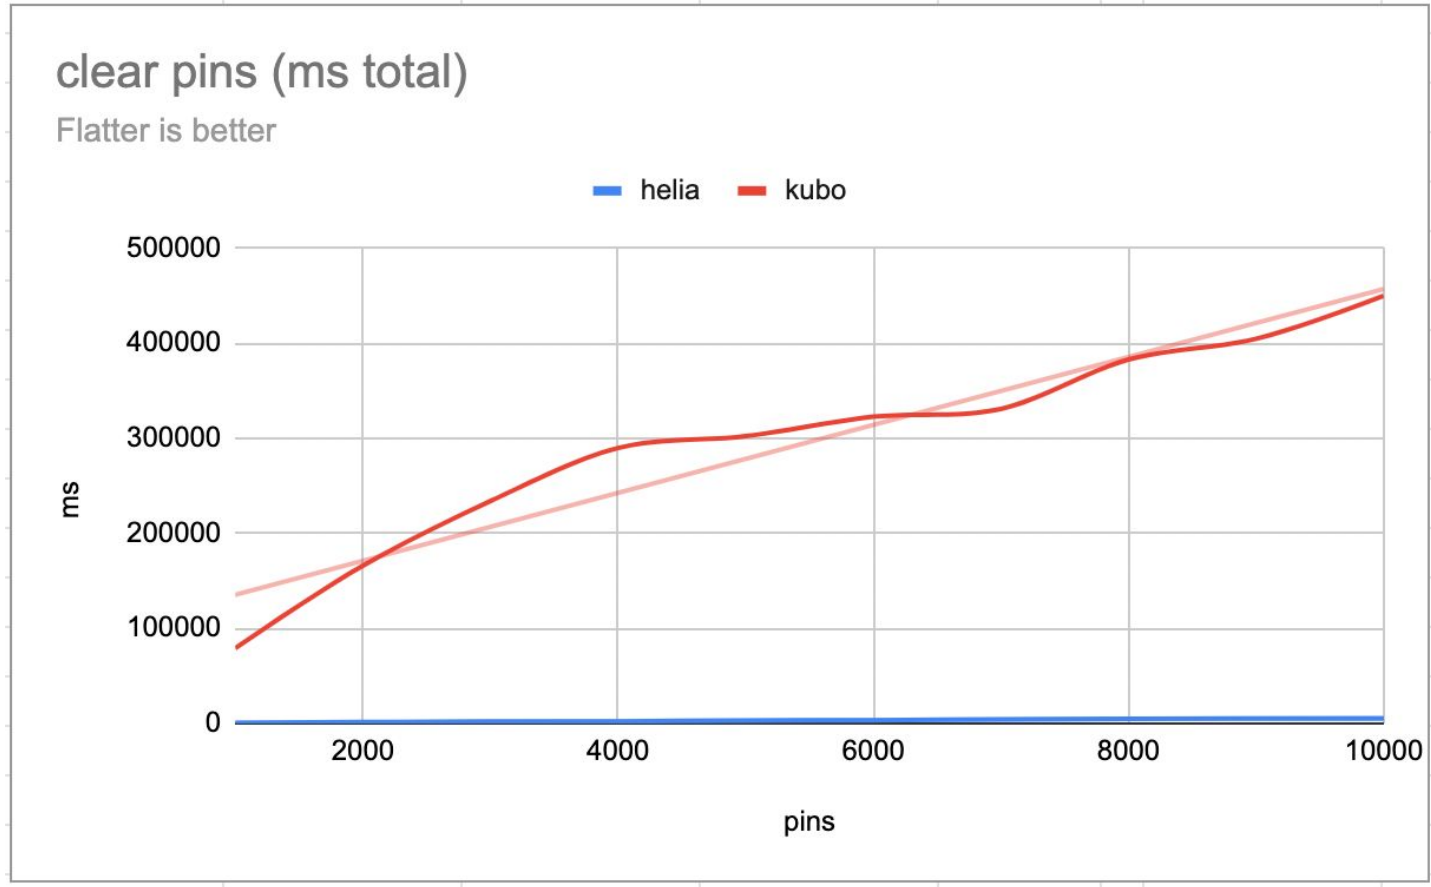




- **Uses reference counting for pins**
- **<https://github.com/ipfs/helia/pull/36>**
- **Benchmark suite is in the helia repo:**
  - <https://github.com/ipfs/helia/tree/main/benchmarks/gc>

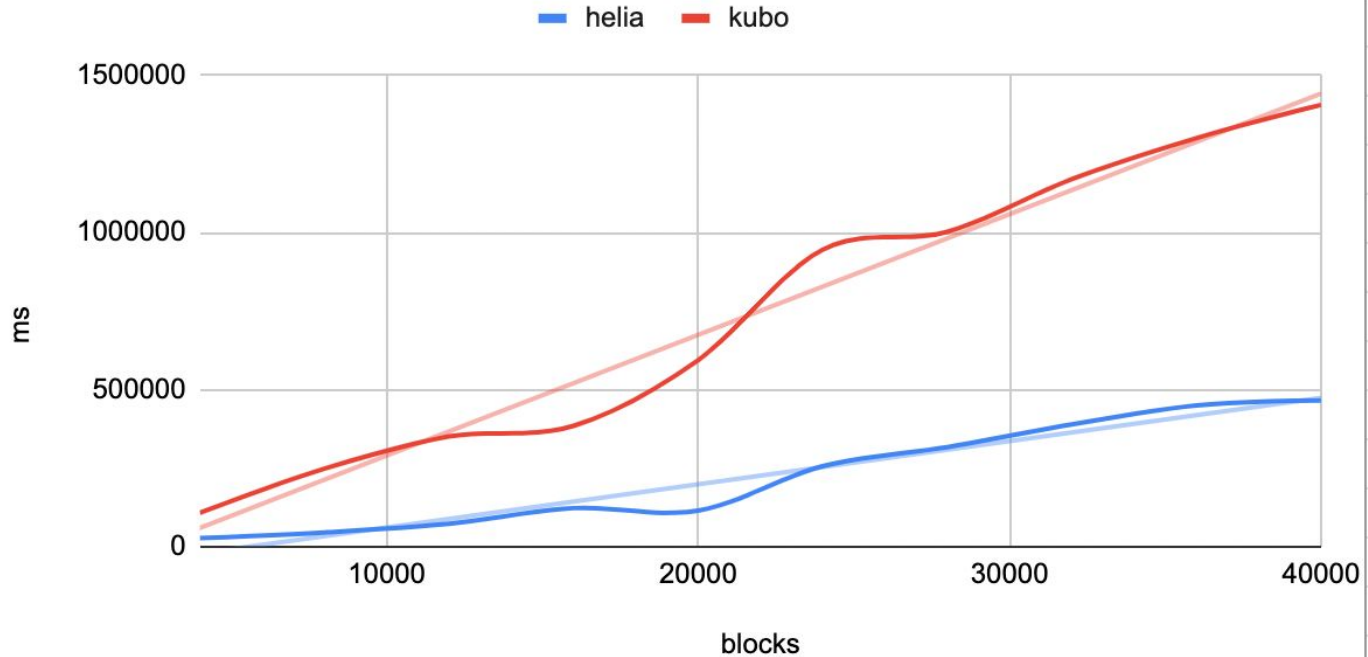


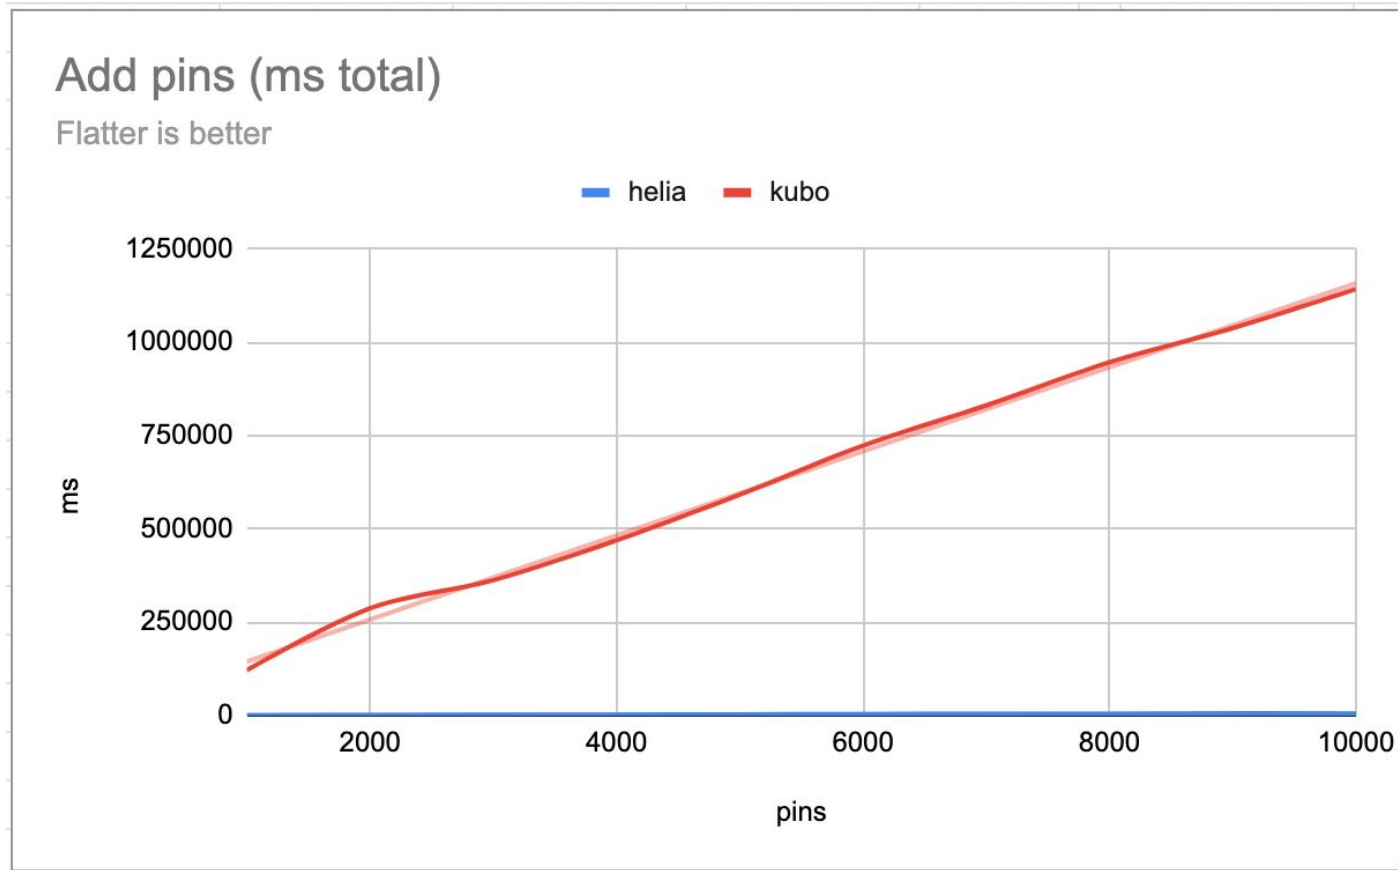




## Add blocks (ms total)

Flatter is better






# DAGs and Bitswap






## What is a DAG?



## What is a DAG?



IPLD Dag



**...maybe don't?**





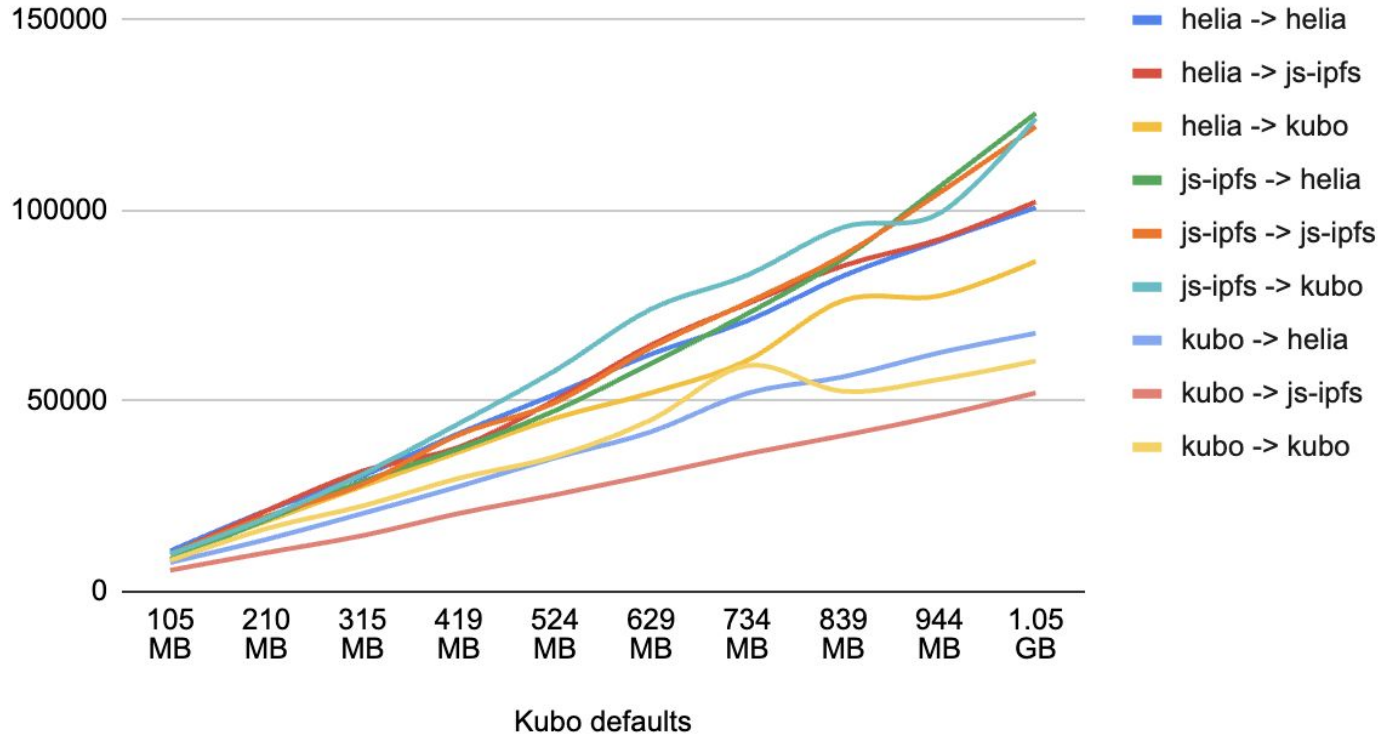
- **Benchmark suite is in the helia repo:**
  - <https://github.com/ipfs/helia/pull/90>





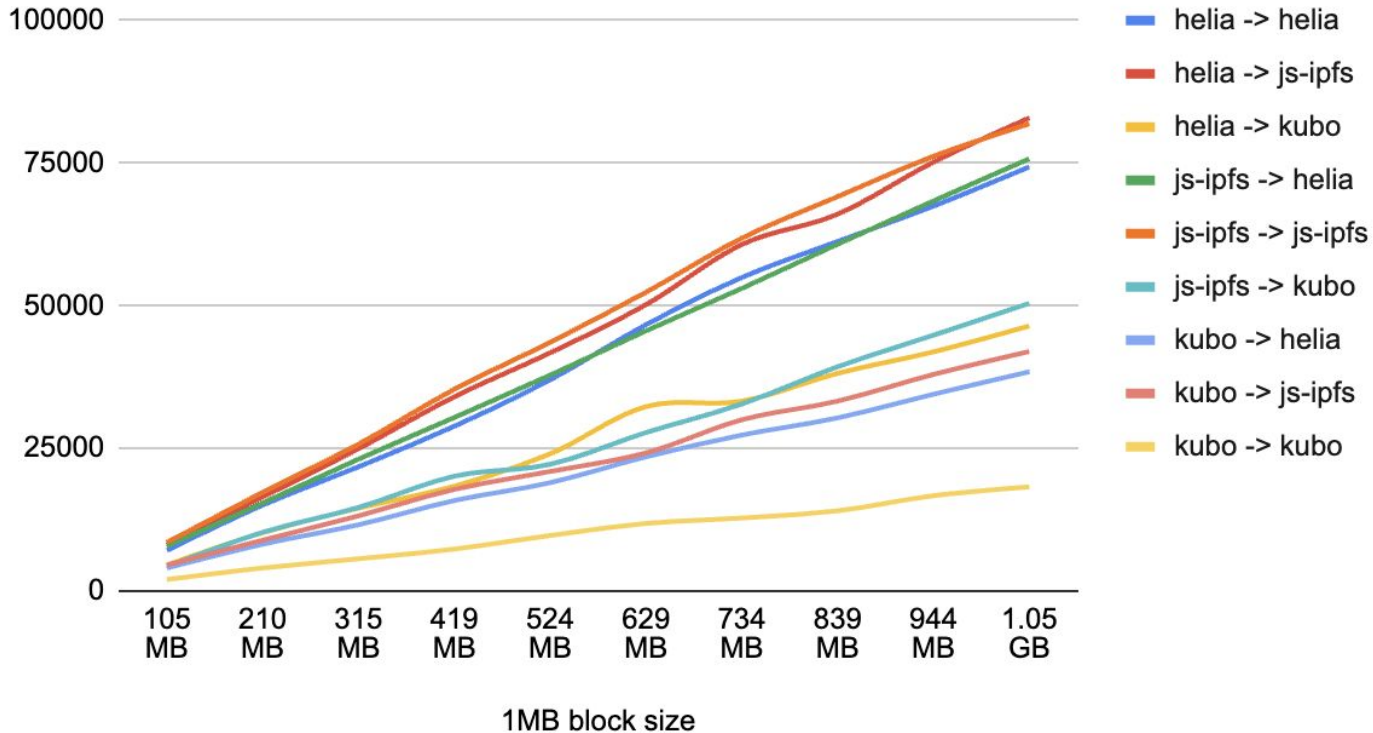


## Kubo defaults





## 1MB block size






## Helia -> Helia

	<b>256 KiB</b>	<b>1 MiB</b>
<b>105 MB</b>	10683ms	7111ms
<b>1.05 GB</b>	100790ms	75353ms

## Kubo -> Kubo

	<b>256 KiB</b>	<b>1 MiB</b>
<b>105 MB</b>	8184ms	2012ms
<b>1.05 GB</b>	60525ms	18262ms





**The end!**

- <https://github.com/ipfs/helia>
- <https://github.com/ipfs-examples/helia-examples>

**Alex Potsides**  
**@achingbrain**