

Beyond Swapping Bits

Accelerating file-sharing in P2P networks and IPFS with Bitswap

Alfonso de la Rocha

Research Engineer

ResNetLab (alfonso.rocha@protocol.ai)

 @adlrocha

 adlrocha.substack.com



Protocol Labs
Research



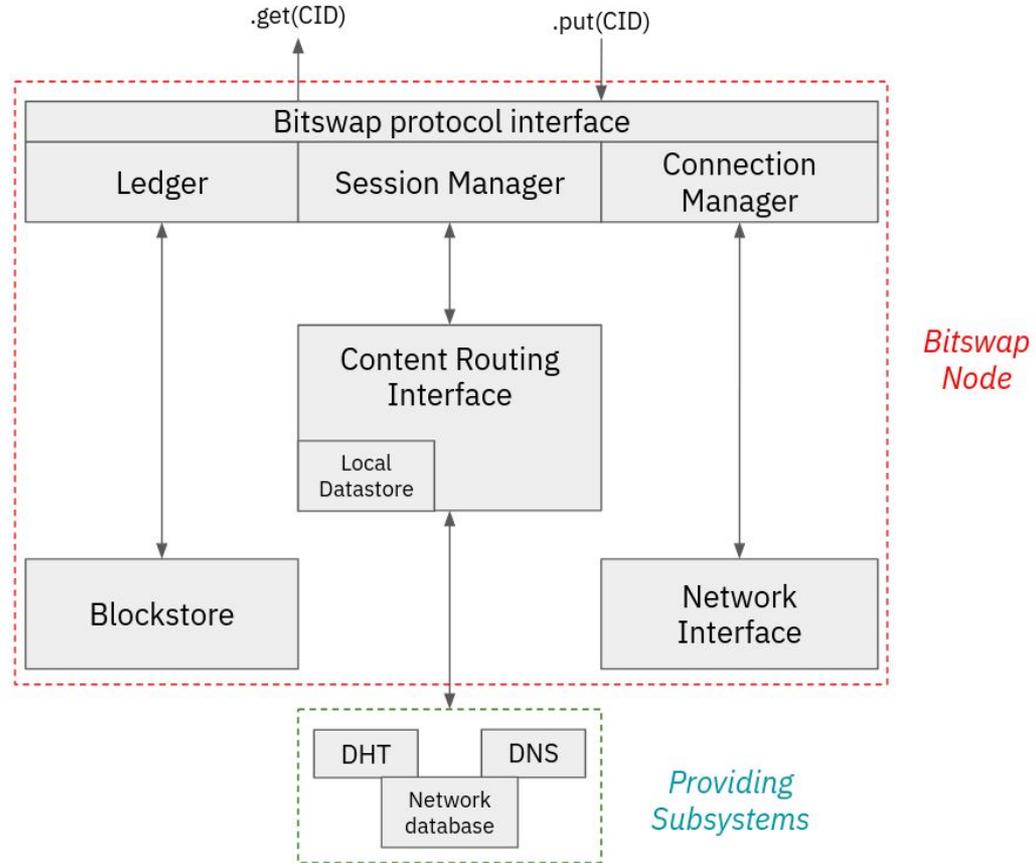
File exchange in P2P networks is hard! 🤖

- Content discovery, resolution and delivery.
- Without any central point of coordination.
- A gamut of content routing systems helping in this quest:
 - Bittorrent: Trackers
 - Web 2.0: DNS
 - P2P networks: DHT (slow 🐢)

Bitswap

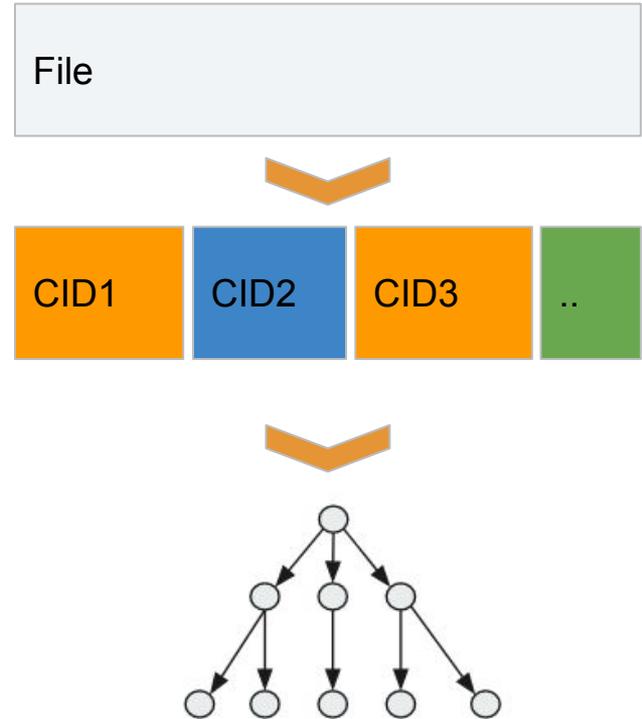
Bitswap message-oriented protocol that helps content routing subsystems to overcome their trade-offs.

- IPFS' exchange interface
- Filecoin's block synchronization



Content in Bitswap

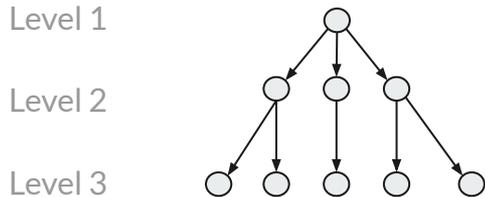
- Content is chunked in blocks
- Blocks are uniquely identified by a **Content Identifier (CID)**, i.e. hash of the block
- Structured as a DAG (link of blocks)



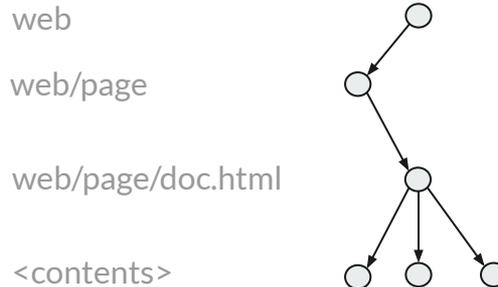
Request Patterns

IPFS requests blocks from Bitwap. Two common request patterns:

Request all nodes at each level of a DAG (eg a file)



Request nodes down a DAG (eg web/page/doc.html)

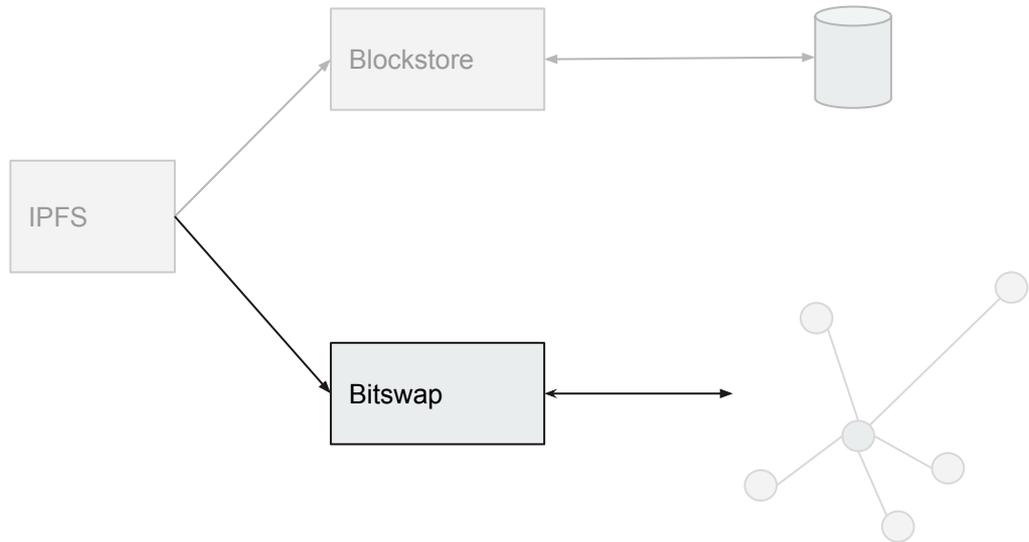


Fetching files

Bitswap is the exchange interface in IPFS

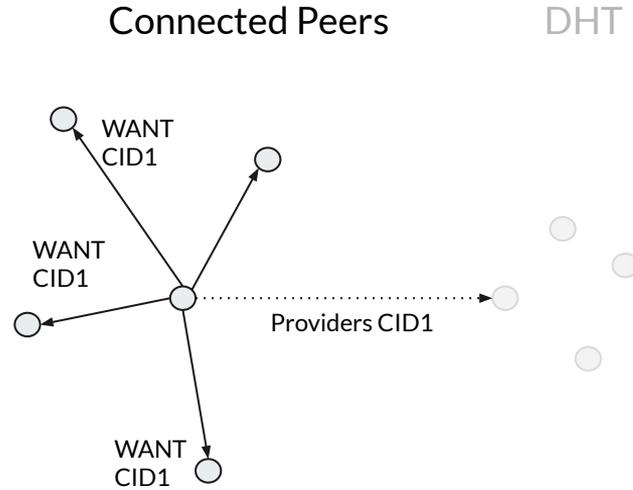
IPFS calls Bitswap to gather files from the network

- Requests: WANT-HAVE / WANT_BLOCK / CANCEL
- Responses: HAVE / BLOCK / DONT_HAVE



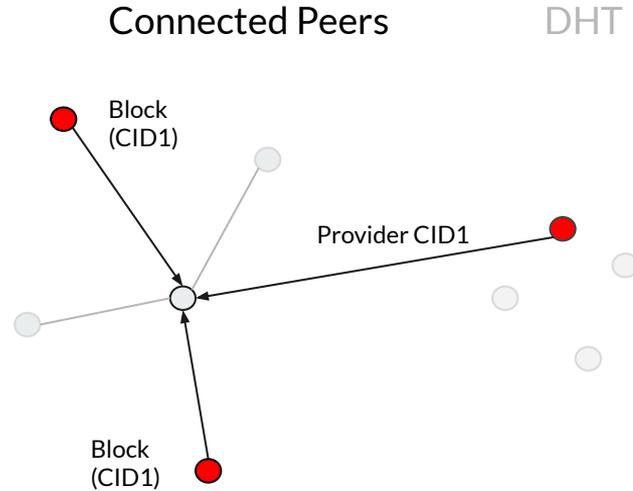
Bitswap - Discovery

- Broadcast WANT to connected Peers
- If there's no response, ask DHT who has root CID



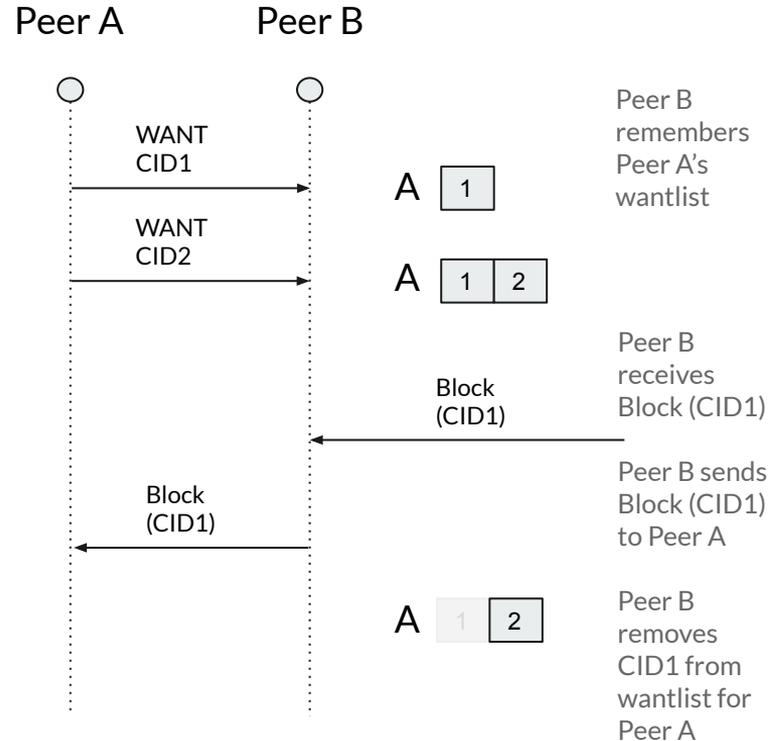
Discovery - Sessions

- Peers who respond are added to the **Session**
- Subsequent requests are sent only to peers in the session



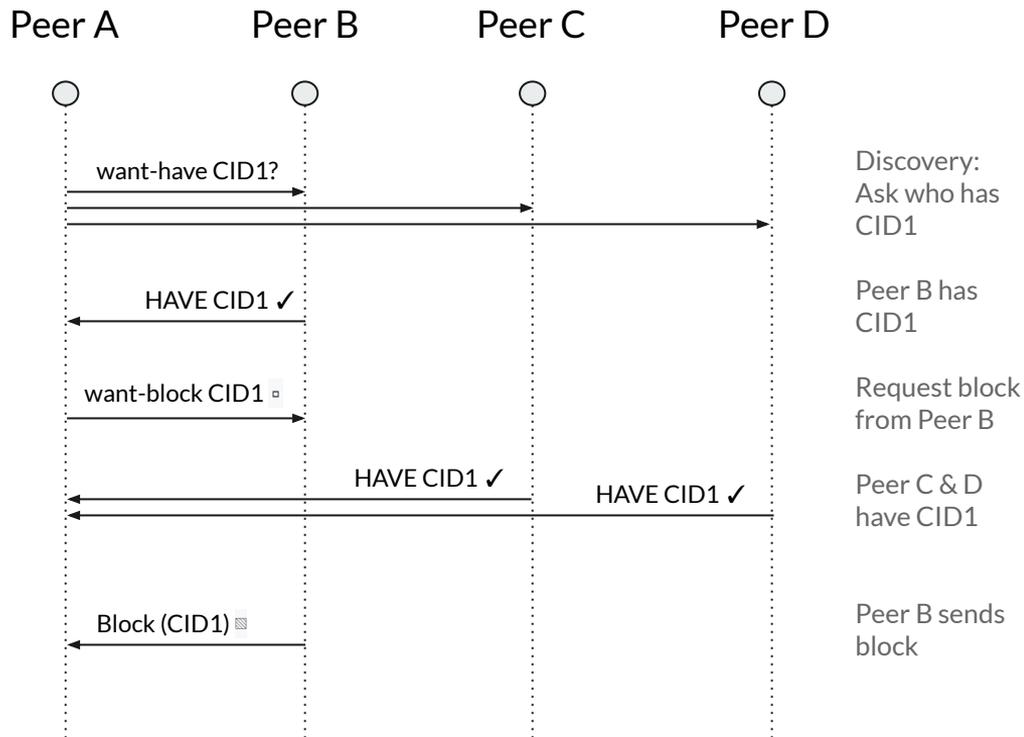
Discovery - Wantlists

- Nodes send WANT messages to peers
- Each node **remembers the want list** for each of its peers
- The wantlist is discarded when the peer disconnects



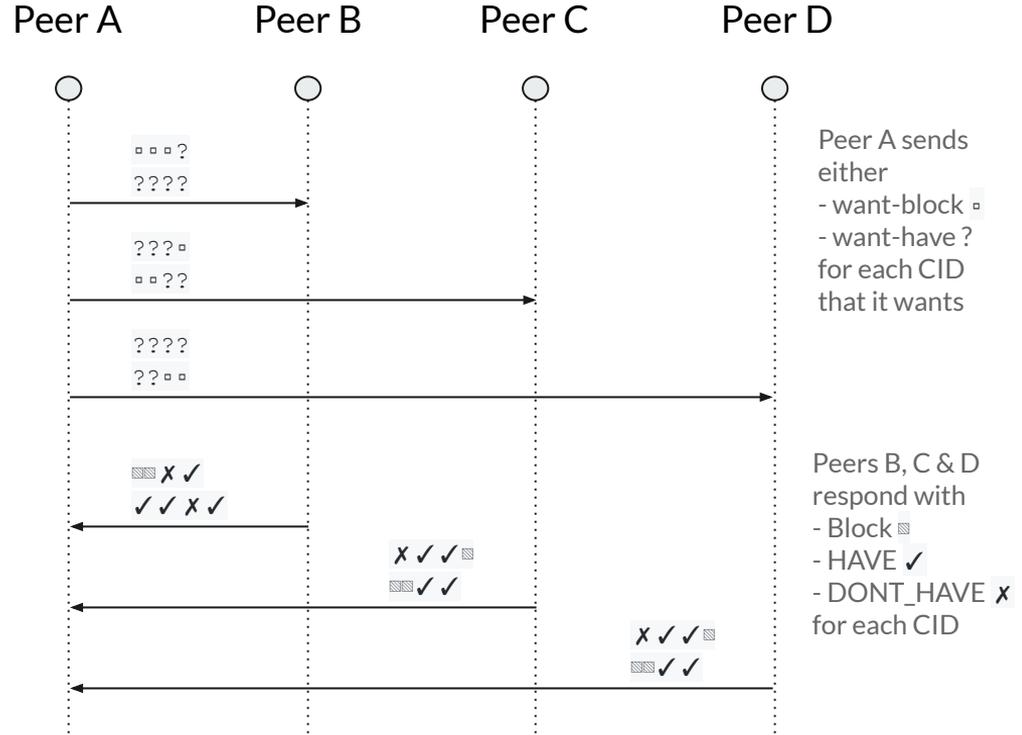
Discovery - Transfer Roundtrip

- **HAVE message**
 - Sometimes we don't want a whole block
 - We just want to know *who* has a block (eg for discovery)
- **Two kinds of WANT message**
 - want-have
 - want-block
- If the block is small enough, send the whole block (instead of sending HAVE)



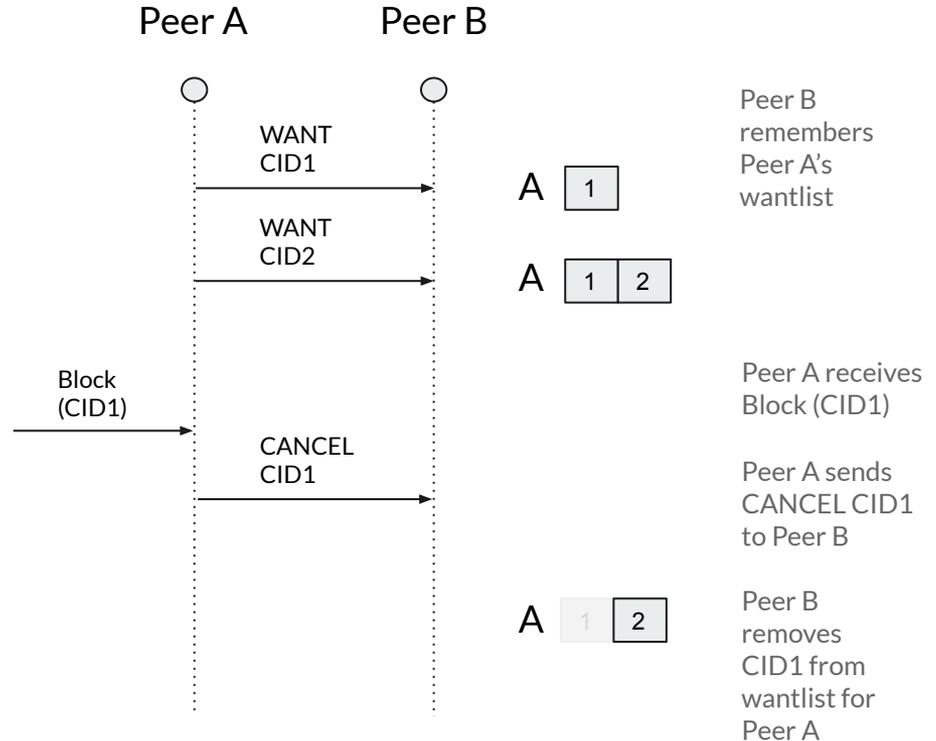
Discovery - WANT-HAVE BROADCAST

- **DONT_HAVE message**
 - Allows peer to indicate that it does NOT have a block
 - Requestor can set a flag to tell responder to send DONT_HAVE in response to want-block or want-have
- **Requests:**
 - want-block
 - want-have
- **Respond with combination of**
 - HAVE
 - DONT_HAVE
 - block

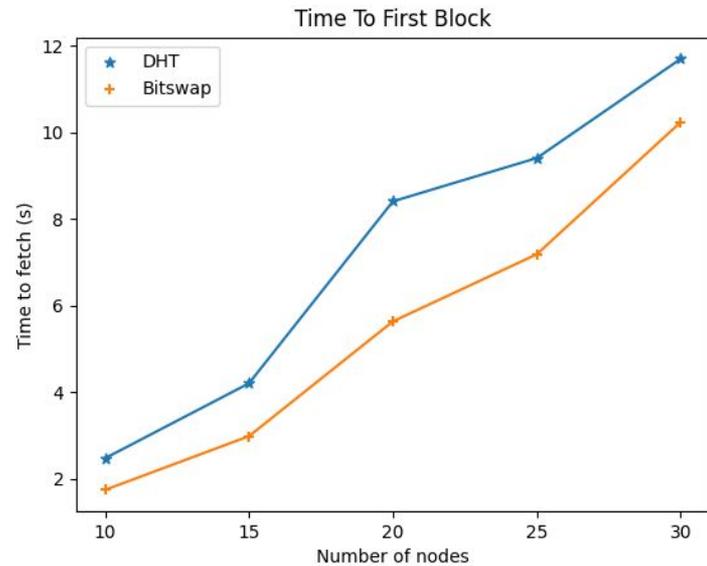
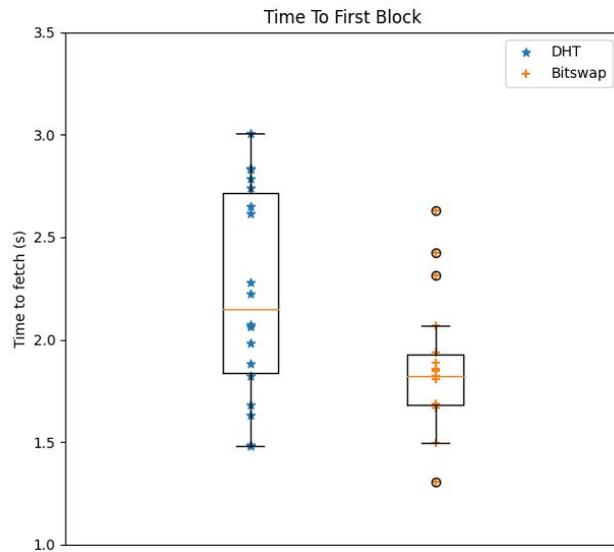


Wantlists - Cancel

- When a node receives a block it wanted, it sends a CANCEL message to all peers it has requested the block from



Bitswap v.s. DHT





Bitwap Issues and Beyond Swapping Bits

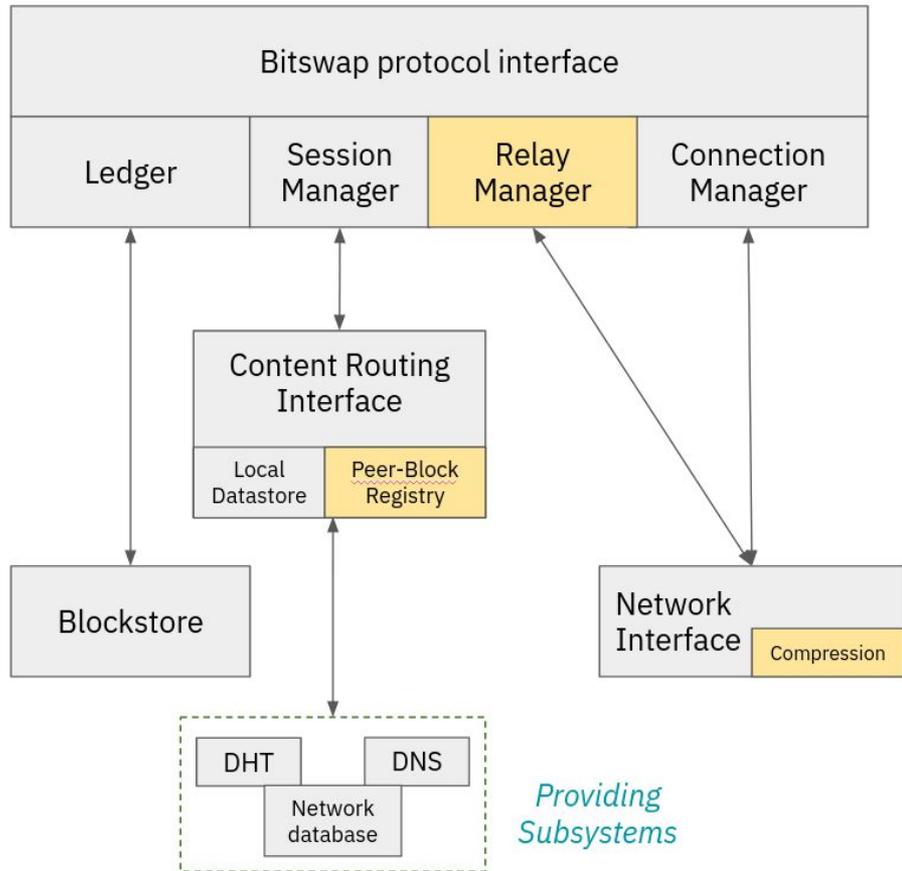


- Current **“one-size-fits all”** implementation may not suit every use case.
 - No way to configure the protocol to fit the client’s needs.
- **Blind and deterministic search.**
 - We don’t use a priori information of the protocol (or other protocols) when we start the discovery.
- **“Dumb” requests:**
 - Split requests and use multiple non-overlapping transfer streams
 - Use selectors and queries to discover full DAG sections and request their transfer.
- More **efficient use of bandwidth.**



Ongoing work: <https://github.com/protocol/beyond-bitwap>

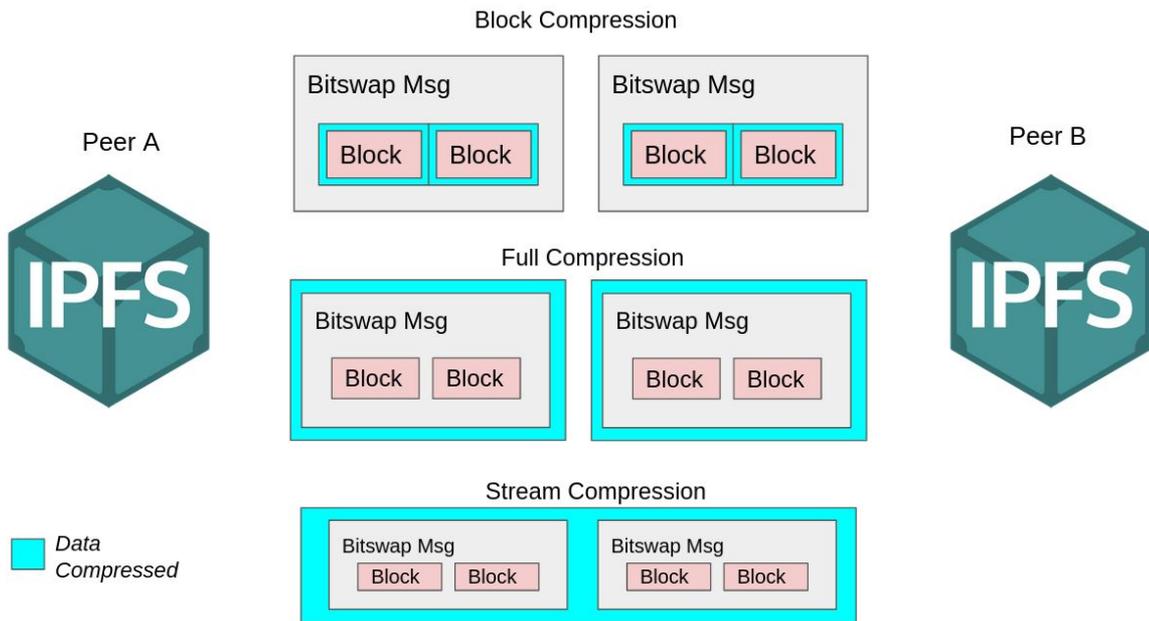
Beyond Swapping Bits



Compression in Bitswap



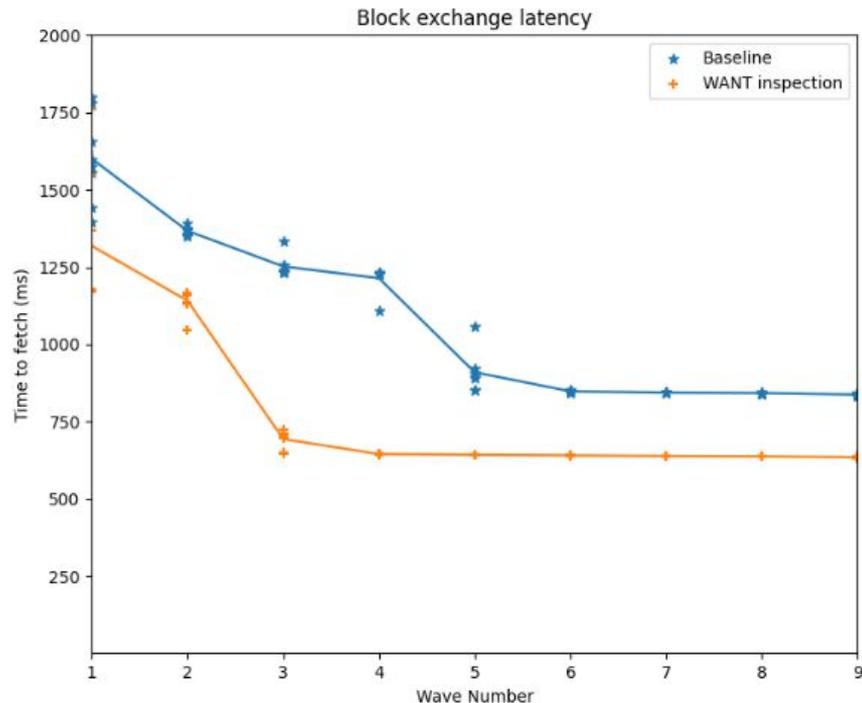
- Up to 75% on bandwidth savings



WANT message inspection

Nodes requesting blocks will potentially have it in the future

Inspect WANT messages received to direct subsequent discoveries for content.



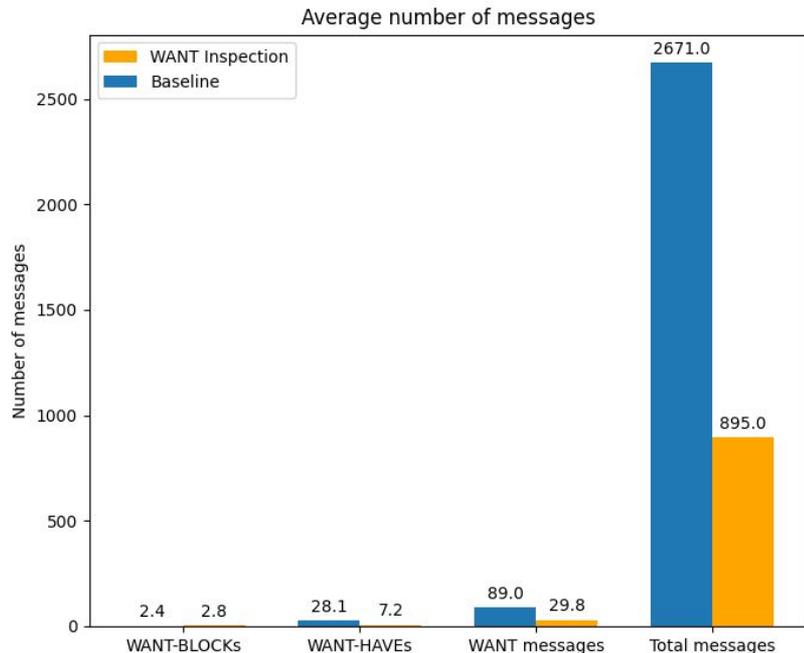
<https://research.protocol.ai/blog/2020/two-ears-one-mouth-how-to-leverage-bitswap-chat-ter-for-faster-transfers/>

WANT message inspection

Nodes requesting blocks will potentially have it in the future

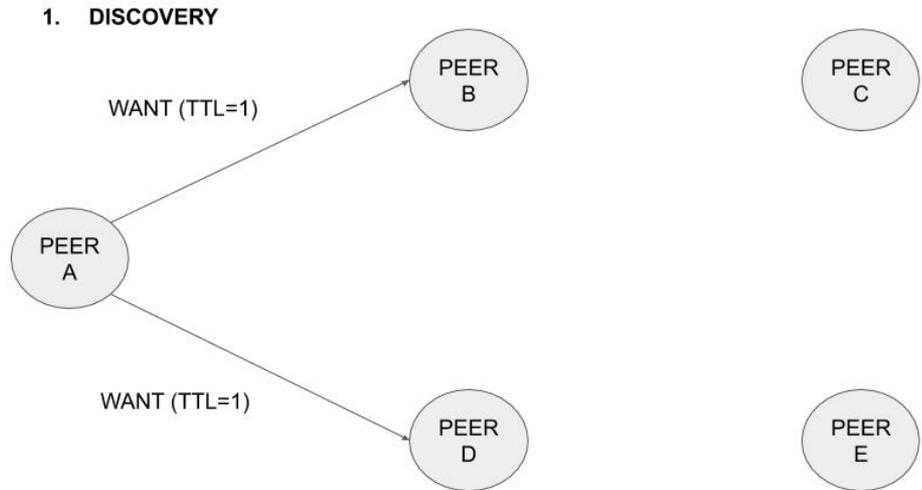
Inspect WANT messages received to direct subsequent discoveries for content.

<https://research.protocol.ai/blog/2020/two-ears-one-mouth-how-to-leverage-bitswap-chat-ter-for-faster-transfers/>

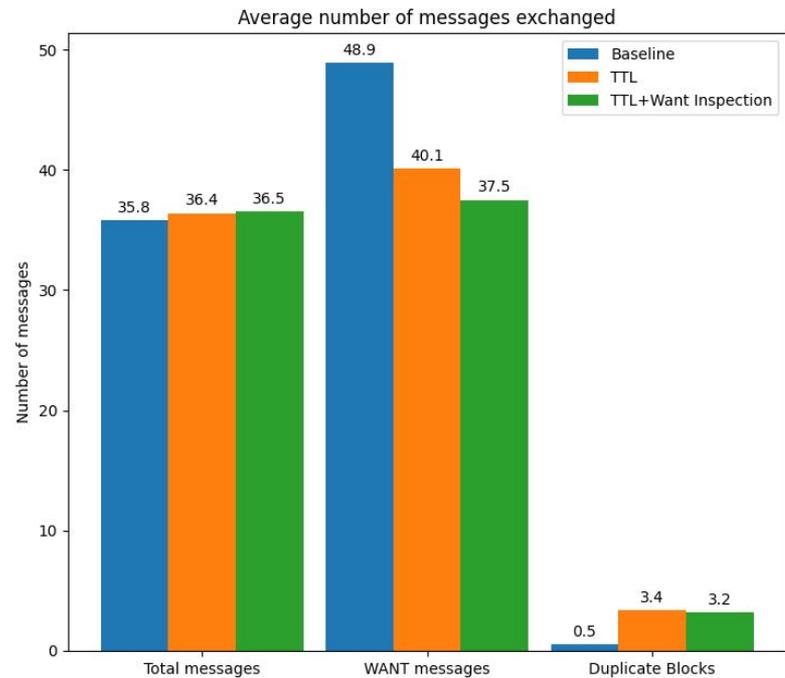
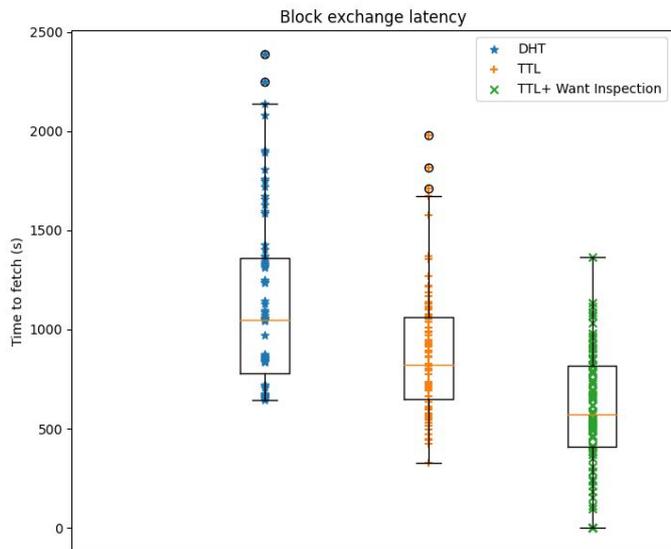


Jumping Bitswap

Increase range of discovery of Bitswap without resorting to providing subsystem.



Juming Bitswap





Ongoing Research

- RFCs with potential improvements.
- Research and development teams building prototypes for the RFC and coming up with new ones.



Ongoing Research

- RFCs with potential improvements.
- Research and development teams building prototypes for the RFC and coming up with new ones.

Help us make file-sharing in P2P networks blazing fast! 

<https://github.com/protocol/beyond-bitswap>

Questions?

We ❤️ Feedback 

Reach out if you want to contribute to the work!
Join the discussion!

Alfonso de la Rocha

Research Engineer

ResNetLab (alfonso.rocha@protocol.ai)



@adlrocha



adlrocha.substack.com



Protocol Labs
Research

Beyond Swapping Bits

Accelerating file-sharing in P2P networks and IPFS with Bitswap

Alfonso de la Rocha

Research Engineer

ResNetLab (alfonso.rocha@protocol.ai)

 @adlrocha

 adlrocha.substack.com



Protocol Labs
Research