Murmurations

Making Movements Visible

Whitepaper v0.1.0 - September 2019

Adam McKenty & Oliver Sylvester-Bradley

Distributed data sharing for the new economy

Contents

Summary

Murmurations is a distributed data sharing network and open source plugin suite intended to help connect new economy projects and organisations and make them visible to the world.

Murmurations includes four components:

1. An extensible data schema for describing organizations and projects
2. An easily-installed plugin and other tools for organizations and projects (nodes) to publish information about themselves to the network
3. An index that maintains a list of all nodes
4. An aggregator plugin that uses Murmurations data to automatically build maps, directories, and feeds that showcase nodes in particular sectors, networks, or regions

Alpha versions of the schema, plugins, and index are being tested online with participating organizations.

This whitepaper describes the background and motivation for Murmurations, its architecture, an initial roadmap for development, adoption, and governance, some risks and associated mitigation strategies. Feedback is requested.

Background

Our current industrial economies are both highly concentrated and proving unsustainable and unhealthy for planetary and human well-being.

Alternatives to mainstream economies are by nature distributed, local, and diverse. Hundreds of thousands of groups, businesses and organizations around the world are working from the ground up to change how economies work and make them more just and sustainable. But these pioneers lack the financial power of major corporations, and can be difficult to see.

Much work has been done to map nodes of social, economic, or ecological change, and make them visible. The typical approach to building a directory or map is centralized: collect information about nodes and add it to a database that is owned and designed by the aggregator. Once established, the aggregators either solicit additions and updates from nodes, manually curate the node records to add new ones and delete or update obsolete ones (or both), or simply allow the database to gradually become out of date.

Several challenges impact mapping work, including:

• Manual effort required to identify and enter the information for many small organisations within the target sector or region makes mapping relatively slow and expensive

• Difficulty sharing data between similar or overlapping maps, directories, or aggregators leads to redundant effort to map the same entities for different purposes

• Requirement for a single node to update data on multiple platforms (each with their own set of fields and authentication requirements) makes it onerous for nodes to keep their own data up to date

• Fragmented, incomplete, out of date data make it more difficult for searchers to find organisations or projects to support and collaborate with, and reduces possibilities for trade within the solidarity economy.

In order for the decentralised economy to flourish, better systems of coordination and collective intelligence are needed. Murmurations aims to contribute to solving some of the problems above using bottom-up distributed data sharing.

The Murmurations approach

Nodes hold their own data

In the Murmurations network, every node holds the authoritative copy of data about itself. This is shared through the network, allowing maps, directories, and other aggregators to stay automatically up-to-date by pulling the authoritative data held by the node itself.

Start as simple and low-tech as possible

Much of the alternative economy does not operate on cutting edge technology. In order for a coordination system to be relevant, adoption needs to be smooth and simple for the majority of the projects and organizations who could be part of the network. Our approach is to start with a suite of plugins for Wordpress, which is used by the worldwide majority of SMEs and small NGOs ([35% of all websites](#)), while also releasing code to generate static Murmurations-compatible structured data files that can be uploaded to any website, so that any website using any CMS or other publishing system can join the Murmurations Network. Using the plugin, a website administrator with only limited technical knowledge can get a node setup and on the network in less than two minutes.

Use existing standards for open data

Murmurations makes use of existing standards for open and structured data. These include [RDF](#) and [JSON-LD](#) structured data specifications, schemas from [Schema.org](#) and the [RIPESS ESSGlobal](#) vocabulary for Social and Solidarity Economy organizations, and [RSS](#) for news and event data. We do not want to reinvent the wheel and aim to build upon existing standards and technologies wherever possible.

Build in flexibility and extensibility

Starting with a minimal specification for node data and a basic plugin architecture, Murmurations can be extended with add-on schemas that specify fields relevant to particular networks, regions, or sectors. Custom node plugins and aggregator code allow the creation of new functionality on top of the Murmurations data structure and data sharing architecture.

Architecture

Schemas

Schemas specify how data is structured in the murmurations network. Schemas are stored in JSON files that are read by the node plugin, and define the attributes of each data field, including how it shows up in the input form and how data is validated once entered.

Base schema

Murmurations uses a simple starting set of fields (based on Schema.org's "Organization" schema and the ESSGlobal Social and Solidarity Economy vocabulary) as the base schema for describing nodes. The current version is housed in the Murmurations schema repository. This base schema is under development, and feedback is requested. Comments can be made via issues on the GitHub repo, or via the contact form on Murmurations.network.

Add-on schemas

Add-on schemas can be used to define fields that are relevant to particular networks, organization types, sectors, or data applications.

In the current Murmurations architecture, network organizations can host add-on schemas, and register them with the murmurations index. Once an add-on schema is registered with the index, nodes that identify themselves as members of the network that owns the schema have the option to fill fields defined by the add-on schema within the node plugin.

Example

The Worker Co-op Alliance of Vancouver wants to show a map of their members, including how many worker-owners are part of each co-op. The Alliance's website administrator creates an add-on schema that defines a field for 'Number of members in the co-op'. Once this add-on schema is published to the index, nodes can choose to identify themselves as part of the Worker Co-op Alliance of Vancouver's network. The new field for 'Number of members', created by the Alliance's admin, automatically appears on the' Murmurations admin pages of nodes which add 'Workers Co-op Alliance of Vancouver' in the "Networks" field, so that these nodes can specify how many members they have, and that data can be used by The Alliance's aggregator.

Specifications for add-on schemas are in the readme document in the Murmurations schema repository.

Nodes

The node plugin, in its initial Wordpress version, automatically populates many of the base schema fields with data from the Wordpress database. This makes the initial node setup very simple. The Wordpress plugin uses a [WP REST API](#) endpoint to provide data to the network.

Organizations or projects not using Wordpress can create a static JSON file and upload it to any web-accessible location, add their node to the index, and participate in the network.

The development version of the node plugin is [available on GitHub](#).

Index

The index is the only centralised component of the murmurations network. Its function is to store URLs to facilitate discovery within the network.

Index of nodes

The node index is used by aggregators to find content. The index stores minimal information about each node, including the node's API address, indexed by the node's URL. Aggregators can query the index API to find nodes matching various parameters.

Index of networks

The networks index maintains a list of aggregators that have defined custom Murmurations add-on schemas. This index is queried to populate the network membership options within the node plugin, and to collect addon schemas to build the node plugin form fields.

Murmurations aggregators

Aggregators query the index, and then node APIs, to generate feeds, maps, and directories based on filter criteria. The Murmurations Aggregator plugin demonstrates how to automatically generate maps and directories from the network.

Custom aggregators can build on the demo aggregator plugin code to provide new views, directories, maps and other uses of the Murmurations data.

Aggregators can publish custom add-on schemas so nodes that are members of their organisations, or networks, or within their sector can publish sector-specific data that is of interest to the aggregator. The "networks" field within the node plugin can also be used by aggregators to limit queries to only nodes within their network.

Ally links

In the current plugin and schema, nodes can add links to allied organizations or projects. This will help provide validation data for nodes which specify reciprocal "Allies" and help keep the network spam-free. As 'Ally links' are populated this will also enable Murmurations to present a "social graph" of new economy organizations, and potentially facilitate index-free crawling of the network.

Questions & challenges

Adoption

How will Murmurations be adopted widely enough to make it effective?

In order for Murmurations to work as intended, wide adoption is necessary. Our plan is to bootstrap adoption by working with specific network organisations and sectors to build adoption within their memberships. This strategy includes:

- Data imports — aggregators can import existing data into the network, allowing them to very rapidly build usable maps and directories from a combination of static data and data from nodes on the network. As adoption progresses, static data hosted by the aggregator is seamlessly displaced by self-hosted data from nodes on the network.

- Network on-boarding — working with specific existing network and mapping organizations to create custom schemas and aggregators and build adoption within their memberships.

- Gamification and referrals — building sharing, peer invites, and other outreach tools into the installation and plugin admin.

- Benefit communication — communicating the publicity advantages to nodes that would benefit from greater visibility.

Content quality and inclusion

How will spam or off-topic content be kept out of the network?

Murmurations is intended as an open network. In keeping with the nature of an inclusive, decentralized economy, there are not strict criteria for who should or shouldn't be part of the network. However, as in any participatory internet platform, mechanisms need to be in place to prevent spam and support quality content.

Four tools will be used to keep Murmurations as free from spam as possible:

- **Ally links** will create a credibility graph of nodes on the network
- **Aggregator flagging** will identify off-topic or spammy nodes or feed content
- **Blacklist lookup** — node URLs and submitting IPs can be checked against third-party spam blacklists
- **Node scores** will be calculated based on the three criteria above and the content of node data. Aggregators will be able to automatically exclude nodes whose score is lower than a threshold specified by the aggregator.

Performance

How will the index handle traffic as the network scales?

The murmurations index will need to support requests from nodes (as new nodes install the plugin or existing nodes update their data) and requests from aggregators.

Load level
Because both traffic sources are intermittent, the load on the index is expected to be relatively light in proportion to the number of nodes on the network. Node update frequency is not expected to exceed an average of one update per month per node. Aggregator update frequency is expected to be on the order of once per week for each aggregator. Assuming a 1:100 ratio of aggregators to nodes on the network, this results in an index request frequency on the order of 0.034 requests per day per node on the network (or about 3,500 requests/day with 100,000 nodes on the network).

Mitigation strategies

Queues
Communication between the index and plugins is generally not highly time sensitive. An initial performance strategy on the development roadmap is to add queues to the index, so that requests from nodes can be processed asynchronously as resources are available on the index server.

Distribution
Eventually, the murmurations index could be distributed, removing the performance bottleneck of hosting on a single server. Ideally, the centralized index could disappear entirely, replaced with index data distributed among nodes and aggregators on the network.

Licensing
Murmurations code is licensed under a [Peer Production License](#). This allows it to be freely used by organizations operating within the social and solidarity economy, but prohibits commercial use by conventional profit-extracting businesses. More permissive licences can be negotiated on a case-by-case basis for aligned organizations.

Funding
Murmurations has been an entirely volunteer, unfunded project to date. To be successful funding will be required to develop, maintain, and promote Murmurations. Three approaches could contribute to funding ongoing work.

Grant funding - Murmurations has the potential to make many other efforts to promote sustainable and regenerative economies more effective, and reduce costs for current mapping efforts. This provides a strong case for funding agencies that support existing mapping projects.

Crowd funding - Small grants through crowd-funding or open source funding platforms like [Open Collective](#) could contribute to Murmurations development, especially as adoption progresses and networks receive value from using Murmurations.

Professional services - Developing and deploying custom aggregators and schemas for organizations who would like to use murmurations data in ways that exceed the capacities of the demo aggregator plugin.

Governance and evolution

Who will own and manage the plugin code and protocol development?

For the initial phase, Murmurations is an informal volunteer collaboration between Photosynthesis Communications and The Open Co-op. As it grows, we plan to widen the circle of participation and governance. The eventual intent is for the design of the canonical plugin code, schema, and index to be democratically controlled, under the stewardship of a non-profit or social benefit co-op.

Conclusion

It's hard to navigate and develop what you can't see. Murmurations is an attempt to help the new economy flourish by making its presence and the connections within it visible. Whether this will be successful or not depends on how effectively we can deliver value to nodes and network organizations to promote adoption, and how effectively we can work with other mapping and linked data efforts to create an open source, interoperable ecosystem of data.

Building an alternative economy is not a simple task when the playing field is tilted in favour of large enterprises and hierarchical economies of scale. If we can find ways to leverage the distributed intelligence of the new economy and the growing number of people who want to support it, we can create a network of solidarity and resilience that has the potential to outperform the profit-driven corporate economy. Our hope is that Murmurations will contribute to this.

Request for feedback

Our intention is to make Murmurations maximally useful for those who are working towards a more just and regenerative economy. To do this, we need feedback on the schema and architecture, recommendations for network or funding partners, and code review. Please get in touch!