


An architecture and platform for developing distributed recommendation algorithms on large-scale social networks

Journal of Information Science
2015, Vol. 41(5) 686–704
© The Author(s) 2015
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0165551515588669
jjs.sagepub.com


Alejandro Corbellini

ISISTAN Research Institute, UNICEN University, Argentina

Cristian Mateos

ISISTAN Research Institute, UNICEN University, Argentina

Daniela Godoy

ISISTAN Research Institute, UNICEN University, Argentina

Alejandro Zunino

ISISTAN Research Institute, UNICEN University, Argentina

Silvia Schiaffino

ISISTAN Research Institute, UNICEN University, Argentina

Abstract

The creation of new and better recommendation algorithms for social networks is currently receiving much attention owing to the increasing need for new tools to assist users. The volume of available social data as well as experimental datasets force recommendation algorithms to scale to many computers. Given that social networks can be modelled as graphs, a distributed graph-oriented support able to exploit computer clusters arises as a necessity. In this work, we propose an architecture, called Lightweight-Massive Graph Processing Architecture, which simplifies the design of graph-based recommendation algorithms on clusters of computers, and a Java implementation for this architecture composed of two parts: Graphly, an API offering operations to access graphs; and jLiME, a framework that supports the distribution of algorithm code and graph data. The motivation behind the creation of this architecture is to allow users to define recommendation algorithms through the API and then customize their execution using job distribution strategies, without modifying the original algorithm. Thus, algorithms can be programmed and evaluated without the burden of thinking about distribution and parallel concerns, while still supporting environment-level tuning of the distributed execution. To validate the proposal, the current implementation of the architecture was tested using a followee recommendation algorithm for Twitter as case study. These experiments illustrate the graph API, quantitatively evaluate different job distribution strategies w.r.t. recommendation time and resource usage, and demonstrate the importance of providing non-invasive tuning for recommendation algorithms.

Keywords

Graph databases; graph-processing frameworks; large-scale processing; recommendation algorithms; social networks; work scheduling

Corresponding author:

Alejandro Corbellini, ISISTAN Research Institute, UNICEN University, Argentina B7001BBO, Argentina.
Email: alejandro.corbellini@isistan.unicen.edu.ar

1. Introduction

In recent years, there has been considerable interest in the design and development of recommendation algorithms to assist users to cope with the exponential growth of online social networks. Most experimental recommendation algorithms, particularly those developed in academia, are implemented as single-machine, single-threaded applications [1–5]. Even the well-known Twitter's Who to Follow algorithm is implemented in a single machine to reduce architectural complexity [6]. These implementations face scalability issues when the size of the underlying social graph increases, as usually happens in today's social networks, populated with millions of users. Likewise, in situations where computing resources are scarce, generating recommendations becomes challenging. An alternative is to migrate the recommendation algorithm and/or the underlying graph database to a distributed and parallel platform that supports the execution of the algorithm on a cluster of computers (a networked array of computers). In this context, the natural choice to handle social data is graph databases and frameworks for processing graph-based algorithms.

Distributed graph storage and processing systems have attracted much interest from companies and academic institutions that handle large amounts of interconnected data such as social or geo-spatial data. For example, Facebook's social network has 801 million daily active users that upload more than 500 TB to Facebook's data stores, which are analysed in real time to harvest usage statistics. As of 2013, Twitter's 100 million users daily generated 500 million tweets per day, which are processed to show trending topics and produce targeted advertising. Moreover, the amount of data stored and queried by this kind of application makes traditional databases and even common NoSQL solutions unfeasible to cope with the needed performance levels. Even acquiring statistically significant samples from real-world graphs, like those formed by Twitter and Facebook social networks, has proven to be challenging [7]. In the literature, the Big Data [8] term was coined to refer to large-scale data management, including its analysis and the corresponding support platforms.

In response to this challenge, some developments in the form of graph-specific databases or frameworks for processing graph algorithms have arisen. From a system design point of view, graph databases operate at the data storage level. One example is graph NoSQL databases [9], which are optimized for storing large-scale graphs. Complementary, graph-processing frameworks provide programming facilities for developing graph-based algorithms. Some examples of these frameworks include Pregel [10], HipG [11], Piccolo [12] and GraphLab [13]. Surprisingly, many of these frameworks [11, 13] do not even cleanly support permanent graph data distributed storage, which prevents them from efficiently supporting the evolving nature of graph data in real-world social applications. One extreme case is HipG, which requires population of the main memory of the computing cluster with the graph data from a single computational node prior to processing. Another shortcoming lies in the definition of tuning code in order to adjust the algorithm execution to the underlying hardware and storage characteristics, which is common on small setups of heterogeneous hardware. Most frameworks (Pregel, HipG, Piccolo) do not address this problem as they focus on algorithmic modelling rather than code distribution. To solve this problem, the developer might decide to drop the graph framework solution and use a pure distributed processing framework, for example, MapReduce [14], BSP (Bulk Synchronous Parallel) [15] or ForkJoin [16] – manually handling the mapping of jobs to computing nodes according to a fixed criteria. However, this approach entangles algorithm code with distributed execution concerns, which is known to cause code maintainability issues [17].

Therefore, in our view, there is a need for a new support that takes advantage of distributed data stores and provides abstractions to (optionally) allow users to exploit low-level tuning execution mechanisms, while additionally offering a simple graph-traversal API tailored to recommendation algorithms. In this work, we present an architecture for distributed graph-processing systems as well as its corresponding implementation, to aid the development of tunable recommendation algorithms. The architecture, called Lightweight-Massive Graph Processing Architecture (LMGPA), divides the responsibilities of a distributed graph-processing system into a number of interchangeable modules. Our implementation, based on LMGPA, comprises two software modules: Graphly and jLiME. The former provides a Graph API and implements a persistent support for stored graphs. The latter is responsible for distributing and processing jobs, plus providing middleware-level services such as node discovery and failure tolerance. Moreover, jLiME supports high-level job execution models (e.g. ForkJoin, BSP and MapReduce) that the upper layers (e.g. Graphly) can use to distribute jobs.

On one hand, when creating a recommendation algorithm, the graph API that hides all of the mechanisms involved in querying the graph is vital for achieving code portability and modifiability. On the other hand, in some scenarios the developer must customize the inner mechanisms of code distribution depending on the algorithm being developed. For example, the developer might want to distribute jobs depending on the data layout (i.e. the way data is divided and replicated through the cluster) or the computational nodes capabilities. In most graph-distributed processing frameworks, this involves adjusting the original algorithm or the data layout to the new requirements, which requires some effort or might be impractical. In response, we propose non-invasive job mapping strategies to customize the way job processing is

distributed but without modifying the developer's original algorithm. In essence, these strategies provide a bridge between the algorithms expressed using the Graphly API and the underlying computing cluster and data layout.

This paper is organized as follows. Section 2 summarizes relevant related work on graph databases and graph-processing frameworks. Section 3 describes the architecture behind Graphly and jLiME. Section 4 provides a description of the underlying distributed processing support. Section 5 introduces Graphly and its API, and provides usage examples. Section 6 describes the case study used to show the importance of customizing job distribution and the experiments of the different strategies used. This case study involves expressing an existing recommendation algorithm [18, 19] using Graphly and executing it with real Twitter data. Finally, Section 7 draws some conclusions and list future work.

2. Related work

The architecture proposed stems from two existing types of systems: graph-processing frameworks and graph databases. Several frameworks that support graph-processing algorithms in clusters of computers can be found in the literature. Many of those efforts are based on the well-known Bulk Synchronous Parallel (BSP) model by Valiant [15], which divides the execution of an algorithm in multiple supersteps synchronized by logical barriers in which message passing takes place. As an example, Pregel [10] is one of the most popular BSP-based computational models for graph processing. It introduces an in-memory, vertex-centric approach for modelling graph algorithms. Moreover, there are several open-source, publicly available implementations based on Pregel, such as Giraph,¹ jPregel,² GoldenOrb³ and Phoebus.⁴ Trinity [20] is another graph-processing framework, created by Microsoft providing vertex-centric computation using BSP, and also graph querying capabilities and message passing between computing nodes. HipG [11] uses the notion of synchronizers to model algorithms and nodes to perform vertex-centric computations. Each synchronizer might asynchronously call many nodes and then use a barrier to wait for node execution finalization. Synchronizers might spawn multiple child synchronizers, effectively building a hierarchy of parallel synchronizers.

GraphX [21] is a BSP-based framework that covers the whole graph-processing pipeline, that is, data acquisition, graph partitioning and subsequent processing. It also provides a table view of graphs that allows the user to perform filter, map and join operations. Piccolo [12] uses a distributed key-value store to keep the data (e.g. user profiles) associated with each vertex. The map is partitioned and processed in parallel by a kernel function that generates mutations to the map. After each iteration the mutations are applied to the distributed map through the use of synchronous barriers, similarly to BSP. Moreover, Distributed GraphLab [13] is based upon an asynchronous vertex-centric computation model, where the data and graph modifications are scheduled using a distributed locking mechanism. Thus, GraphLab avoids the penalty imposed by synchronization barriers from the BSP model.

Graph-processing frameworks usually load data from a persistent store and then construct a distributed in-memory representation of the graph. In many clusters, RAM memory is scarce and imposes a hard limit on the size of the graph to be processed. Additionally, many graph algorithms [6, 19, 22–24] require in-memory auxiliary data structures to handle intermediate results, which further limits the amount of memory that can be assigned to graph data. Driven by this situation, we created a graph database that supports graph traversal operations and provides a model for designing distributed graph algorithms over its data. Graph data is kept in a persistent graph database and algorithm execution can be non-invasively adjusted according to different criteria, such as main memory usage, CPU speed, CPU usage, secondary disk usage and vertex storage location, among others.

In the literature regarding distributed graph databases, there are at least two types of stores: RDF (Resource Description Framework) stores and Property Graphs. RDF stores originated with the Semantic Web [25–28] and most of them have been developed for years. The RDF specification provides the < Subject, Predicate, Object > representation of relationships in a graph of entities. The Subject is the origin vertex, the Predicate is the edge label, and the Object is the target vertex. Usually, a Context is added to express where a triple comes from. In the context of graphs, the Context is the graph where a triple belongs to. The Subject, Predicate, Object, Context (SPOC) representation is usually called a Quad. RDF stores usually persist the SPOC tuples into a key-value store, and create several indexes to support different types of query access. For example, YARS2 is a distributed RDF store that creates a SPOC index, and four additional indexes to support alternative access patterns. Bigdata [27] is based on concepts introduced by YARS2 [29], but adds an inference engine and an implementation of the Gather–Apply–Scatter graph-processing model. Virtuoso [26] is another example of a mature database that supports the RDF format and provides an inference mechanism over its tuples.

Property graphs are a type of graph store where edges and vertices are modelled with an arbitrary number of properties, usually, in a key-value format. A common technique to store a property graph is to separate the different elements of a graph: vertex adjacency lists (vertices or edges), vertex data, edge input vertices, edge output vertices and edge data. Implementations of this type of databases are fairly new and are mostly related to the NoSQL movement. As an example, the Titan [30] graph database lets the user select the actual database storage support from a number of lower-level data

stores, including HBase [31], Cassandra [32], BerkeleyDB [33] and Hazelcast [34]. Similarly to many new databases, Titan implements the Blueprints [35] interface, which connects it to a number of already implemented tools for graph querying and processing. Other databases provide a graph interface over a Document-oriented store (i.e. a key-value store where values have a known but flexible format, such as XML or JSON), which is the case for OrientDB [36] and ArangoDB [37]. As a final example, FlockDB [38] is an adjacency graph store that, as its name indicates, only stores the structure of the graph. It was developed at Twitter to store the followee/follower relationships of its users. All in all, our structural graph database is based on a simple distributed key-value store, where each key represents a vertex ID, and each value contains the associated vertex data. The vertex ID is used to distribute the values across the computing nodes in a cluster. Then, a Graph API is built over the key-value store to provide graph operations, for example, obtaining the outcoming and incoming edges of a node.

In this work, we propose an architecture for graph storage and processing that, among other features, provides an API to implement traversal algorithms, that is, algorithms that walk through the graph and might aggregate results, generating sums or counts of elements. Unlike the APIs proposed by other frameworks and databases, we focus on the problem of customizing the distribution of the graph algorithm on heterogeneous clusters. This involves providing access to computing nodes and cluster information such as memory usage, CPU usage and job queues size, and access to information about the storage support such as data location (e.g. in which node an adjacency list resides) and cached data location. Thus, using this information, we provide predefined job mapping strategies to increase algorithm performance. Moreover, an extension mechanism is provided for further customization through ad-hoc strategies.

Indeed, job scheduling or job mapping is a well-studied problem in distributed computing [39]. On heterogeneous clusters, that is, clusters of computers with different capabilities (main memory, CPU speed, number of CPU cores, storage capacity), the allocation of jobs to computing nodes allows distributed and parallel applications to maximize a given performance criterion [40]. In our case, when the recommendation algorithm queries a set of vertices, a user-defined mapping strategy groups those vertices into tasks and maps each task to a set of computing nodes. As an example, a locality-aware mapping strategy might group vertices by their location in the storage support.

The proposed architecture is intended to ease the development of recommendation algorithms over social networks, particularly those requiring complex navigation patterns of the social graph. This category includes link prediction problems, which compute which links are more likely to appear in the future on the graph [41]. Depending on the type of graph, such prediction might include friends, followees, products, co-workers, researchers and Web pages. There are several algorithms that use the structural information of the graph to make a prediction. Typically, these algorithms compute the similarity between nodes based on their neighbourhoods or ensembles of paths using local or global indices [42–44]. More advanced algorithms involve complex graph traversal operations. For example, PageRank [23] is a well-known algorithm that ranks Web pages on the entire Web hypergraph using the ‘random surfer’ notion, that is, a Web user that clicks randomly on HTML links of arbitrary Web pages. Hence, PageRank ranks Web pages by the probability that a user will be at a given Web page at any time. Personalized PageRank is a variation of the original algorithm that computes rankings based on a user’s preferences. Similarly, HITS [22] assigns an ‘authority’ value and a ‘hub’ value to each vertex of the graph using the number of outgoing edges and the number of incoming edges, respectively. HITS was originally used to rank Web pages resulting from a Web search, but later adapted to perform recommendations. SALSA [24] is an algorithm that uses the ‘hub’ and ‘authority’ notion of HITS, and the ‘random surfer’ model from PageRank. Both SALSA and HITS receive an initial group of results from a topic search to build a ranking. More recently, the Who to Follow [6] user recommendation algorithm developed by Twitter uses Personalized PageRank to build the initial group of results, and then uses that group as an input to a SALSA-based recommendation algorithm. Other examples of followee and social recommendation algorithms in the same line are described in the literature [45–51]. This type of algorithm relies on an efficient exploration of the underlying graph for generating recommendations and, consequently, solutions optimized for this domain become critical.

3. Lightweight massive graph processing architecture

In this section we describe LMGPA, a reference architecture for developing systems supporting efficient graph-traversal algorithms over large-scale graphs. Examples of such systems include not only social recommendation systems (recommendations on large social networks), but also geolocation-based systems (supporting transportation systems on large cities), bioinformatics systems (processing of large DNA sequence graphs) and security audit systems (discovering attack patterns on networking logs), among others. Moreover, the proposed architecture focuses on systems that require an easy deployment and fast user adoption.

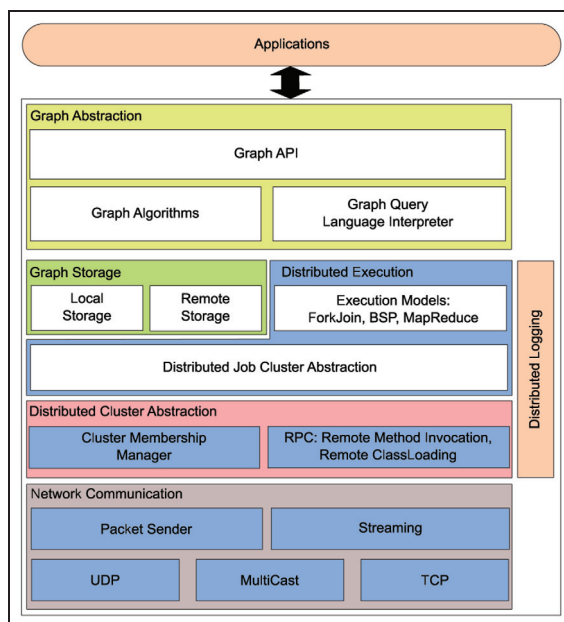


Figure 1. LMGPA layered view.

LMGPA, depicted in Figure 1, follows the Layers pattern [52] in which each layer exposes a set of operations through an API and relies on the API of the lower layer to implement those operations. User applications communicate with the Graph Abstraction layer, the topmost layer of the architecture.

The Graph Abstraction layer defines the API to be used by the applications. This layer might be used through a declarative Graph Query Language (GQL) or a set of functions implemented in a given imperative language. Examples of GQL are Gremlin [53], Graph QL [54], and SPARQL [55]. On the other hand, as an example of standard graph APIs for imperative languages, the Blueprints [35] specification for Java provides a set of conventions to access a graph, similarly to a JDBC (Java Database Connectivity) for graphs. Additionally, this layer can implement a set of basic graph algorithms like breadth-first search, depth-first search, Floyd's algorithm and Dijkstra's algorithm. Moreover, this layer can implement primitives that underlie many recommendation algorithms such as random walks and circle of trust [6].

The second layer is the Graph Storage layer, which provides temporal or persistent storage support for the Graph API. It may be implemented using a purely graph-oriented storage like Neo4J [56] or it can provide a way to store the graph structure in a less direct way like a relational database or a key-value store. If preferred, the distributed support for this layer may be implemented using the Distributed Execution layer explained below.

The third layer corresponds to the Distributed Execution layer. This layer hides concerns related to distributed programming and provides an API to help the Graph layer and the Storage layer to distribute data and code across a cluster of computational nodes. The Distributed Execution layer provides an API to remotely execute any code using the Job abstraction. A Job is an object that can be serialized (represented as a plain stream of bytes) and executed remotely on other node. To support this abstraction, the Distributed Execution layer must implement Job managing capabilities which include, for example, job synchronization, data sharing, failure clean-up and asynchronous/synchronous responses. Additionally, it can provide simple execution model implementations like fork-join or more advanced models like MapReduce or BSP.

The fourth layer corresponds to the Distributed Cluster Abstraction layer, which provides a view of the underlying networked LMGPA processes as a set of Peers that belong to the same cluster. To this end, this layer implements a cluster membership management protocol and a communication mechanism for Peers using RPC (Remote Procedure Call) which, in turn, involves handling data marshalling and unmarshalling. To implement such support, there are many existing mature RPC frameworks that can be employed, such as Protocol Buffers [57], Apache Thrift [58], RMI [59] and RabbitMQ [60].

Finally, the lowest layer is network communication, which implements networking functionality for sending data and code to remote processes. It must provide packet sending through TCP, UDP and UDP Multicast as well as data streaming. Among other things, this layer must support interface disconnections and network address changes, manage TCP

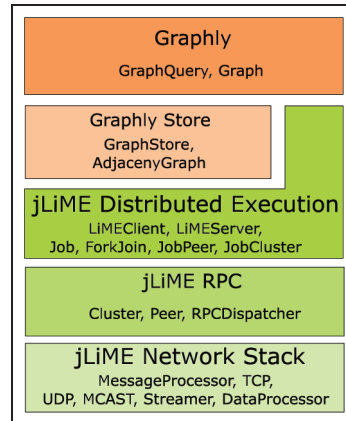


Figure 2. jLiME: a materialization of LMGPA.

connections and identify nodes by ID rather than using a hardware address. Streaming capabilities are crucial for applications where the information being transmitted cannot be handled as a whole in physical memory but, instead, can be processed by chunks in an ‘on demand’ fashion. The network communication layer may be already implemented in the RPC framework used.

4. jLiME: a lightweight distributed execution framework for Java

The LMGPA architecture provides the basis for our own support for graph-based processing applications. This support is implemented via two main modules as shown in Figure 2: the Graphly module and the jLiME module. Graphly implements the Graph Abstraction layer and the storage support for representing graphs. jLiME implements the Distributed Execution layer, the Distributed Cluster Abstraction layer and the Network Communication layer. In this section we will introduce jLiME from a bottom-up perspective.

4.1 jLiME network stack

The bottom part of the framework is the jLiME Network Stack, a reusable and flexible network communication framework, based upon the JGroups library [61]. JGroups lets the developer build a stack of modules that process incoming and outgoing network packets (e.g. modules that detect failures, acknowledge received messages, send and receive packets using UDP or TCP). The jLiME Network Stack not only allows the creation of stacks of packet processing modules, but also allows streaming data between Java processes. In this context, packets are arrays of bytes that are sent over the network as a whole. Streams, however, do not have a fixed size, allowing a sender to send arbitrary amounts of bytes and the receiver to consume the bytes as they arrive. In Appendix A this layer is further discussed.

4.2. jLiME RPC

jLiME provides a Remote Procedure Call module to easily perform remote calls to objects, greatly simplifying the design of distributed programs. First, it introduces the concept of Peer. A Peer is a jLiME process that was detected using any discovery method provided by the network layer. Then, Peers are grouped in a cluster object that allows the developer to select them and keep track of their state.

An RPCDispatcher object allows a method on a given Peer and a given registered object to be called. jLiME also allows the creation, for a given object interface, of a set of helper classes to call methods from that interface as if they were local invocations. One of the helper classes is a factory class for creating a proxy object for a given Peer and a given RPCDispatcher. The usual workflow for using jLiME RPC is shown in Figure 3.

Our Distributed Execution module (Section 4.3) is based on the RPC module. In this case, the RPC module simplifies sending and receiving Job objects to be executed. The Distributed Execution module registers a JobExecutor object on its RPCDispatcher, and then calls the ‘execute’ method on remote JobExecutors using a proxy object, previously generated using the RPC module. Thus, the JobExecutor code contains a minimum amount of code related to remote execution because remote method invocations look like local invocations.

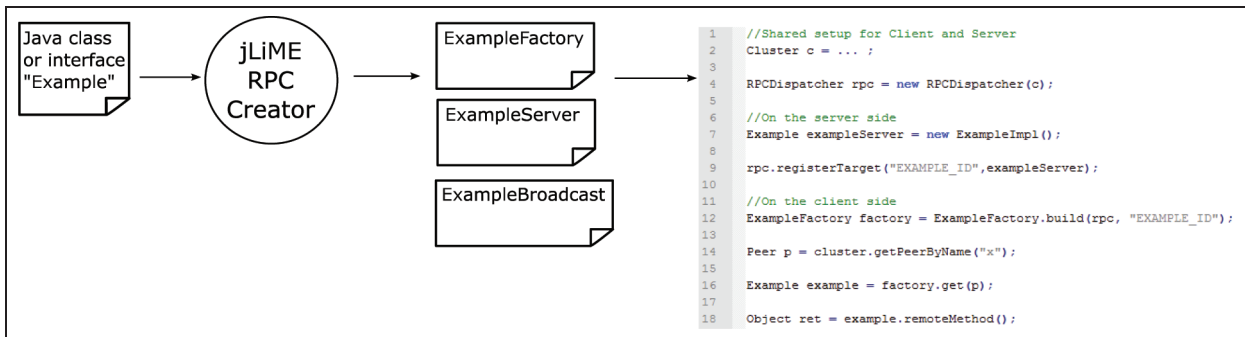


Figure 3. jLiME RPC workflow.

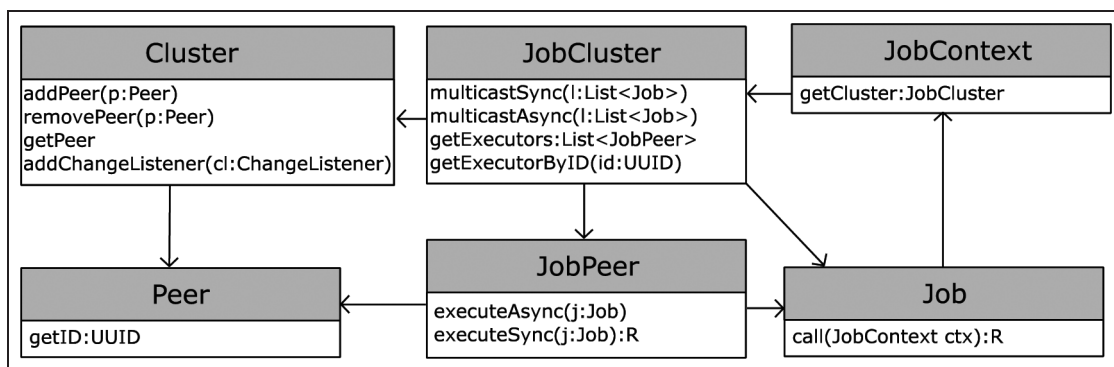


Figure 4. jLiME design: main abstractions.

4.3. jLiME Distributed Execution

The jLiME Distributed Execution module implements the interfaces shown in Figure 4, needed to create a remote job execution environment. The first abstraction introduced by this module is the Job class, which allows a developer to execute arbitrary code on a remote machine. In order for jLiME to run a given Job on a Peer, the developer must obtain a JobPeer object, which encapsulates a Peer instance, and call the execute Job method. A JobCluster object encapsulates a Cluster object and provides access to JobPeer objects. Besides keeping track of JobPeers, the JobCluster divides Peers into job executors, that allow the execution of Job objects, and job clients, that submit Jobs to job executors.

5. Graphly: a simple Java API for graph structure traversal

The graph support developed in this work provides a simple interface to traverse the structure of a graph. A graph *traversal* is an operation that follows the links of a graph, gathering information about links or vertices. We support this operation by relying on two modules: jLiME and the Graphly Store. jLiME helps to distribute any calculation needed by the Graphly API. The latter provides a distributed persistent store for the graph structure. In the following sections we will give an insight into the implementation of the Graphly Store and the Graphly API.

5.1. Graphly Store

Graphly's Store persists the graph data on a cluster of computing nodes. The implementation proposed in this work uses a distributed key-value store to save the adjacency lists of the graph, that is, every vertex is represented by an ID, which is used as a key, and has an associated value comprising a list of incoming or outgoing vertex IDs. Then, a modulo-based hash function to the vertex ID is used to distribute the adjacency lists among computing nodes. Each computing node is responsible for storing a group of adjacency lists on its local database. In our implementation we used LevelDB

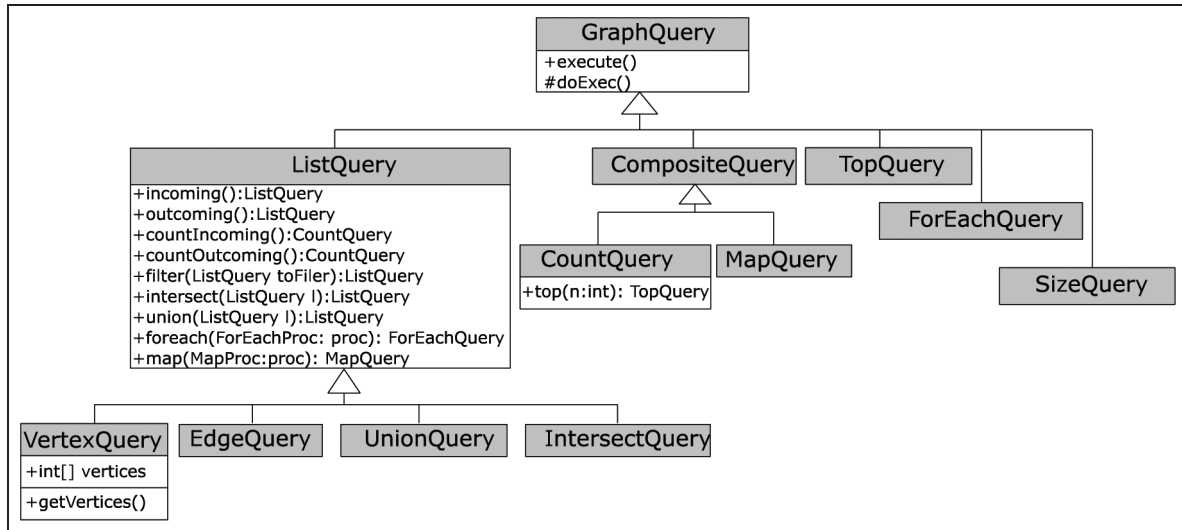


Figure 5. GraphQuery hierarchy.

[62] as the local database implementation. We chose LevelDB because it is a key-value store, which matches our representation of the data, and provides very fast random and sequential access.

5.2. Graphly Traverse API

Graphly provides a querying API for accessing the structure of a graph and exploring it [63]. The class diagram representing the different graph queries is shown in Figure 5. The GraphQuery class provides a common interface for all graph queries.

The most important type of query in our design is the ListQuery, which represents a query that gives as a result a list of vertices (e.g. integer IDs). As a consequence, the most trivial query is the VertexQuery class, which already contains a list of vertices and, when executed, only returns the contained list. This class is used to create more powerful queries and maintain homogeneity with the rest of the GraphQuery hierarchy. The UnionQuery and IntersectQuery classes perform a union or an intersection of executing two ListQuery queries. Finally, the EdgeQuery obtains the list of vertices that has incoming or outgoing edges to the currently contained ListQuery from the Graphly Store. As can be seen in the figure, the ListQuery abstract class provides methods to create several types of queries from a given ListQuery object, allowing the user to chain different types of ListQuery.

Other types of queries perform different operations on ListQueries. The SizeQuery obtains the size of a result from a ListQuery, instead of transmitting the list of elements to the client. The CountQuery counts the number of appearances of a given vertex in a set of results. The TopQuery obtains the top N elements from the result of a CountQuery. The ForEachQuery provides an extension point for a user to create custom queries to be executed on the result of a ListQuery. Similarly, the MapQuery provides a ForkJoin-like mechanism to distribute a given query along the cluster, thus balancing the processing load of the query.

A developer may define different graph traversals algorithms using the interfaces described above. As an example, the Common Neighbours metric [41] was implemented using Graphly's API. Common Neighbours is one of the simplest metrics to measure the similarity of two vertices on a graph and it is a common strategy for link prediction in recommendation systems. For two vertices x and y , let $\Gamma(x)$ denote the set of neighbours of x , their common neighbourhood is defined as $|\Gamma(x) \cap \Gamma(y)|$. Graphly's version of Common Neighbours is:

```

int common_neighbours = graph.vertex(x).neighbours().
    intersect(graph.vertex(y).neighbours()).size();
  
```

The *graph* variable is an instance of Graphly. Using the *getVertex* method, the user obtains a VertexQuery instance for the vertex x . The *neighbours* method creates an EdgeQuery that obtains the vertices having incoming or outgoing edges to x from the Graphly Store. Finally, the result is intersected with the neighbours of vertex y , and its size is returned.

5.3. Graphly job mapping strategies

Before each traversal step, it is possible to decide where to perform the merge of subresults of such traversal. This decision depends on the user and the nature of the algorithm being developed. As an example, if the cluster consists of an heterogeneous group of computing nodes (i.e. different types of CPU, amount of RAM memory, disk storage capacity, etc.), the user may decide to distribute the merging of subresults according to each node's capabilities to achieve lower recommendation times or reducing resource usage (e.g. reducing network usage may be useful in Cloud Computing environments where inter-node communication is charged). Moreover, a CPU-bound algorithm may need to distribute tasks according to CPU capabilities, whereas a memory-intensive algorithm may distribute tasks depending on the amount of RAM memory on each node.

In any case, the algorithm original code remains unchanged. Thus, the strategies effectively work as a transparent connection between the graph algorithm and the underlying platform and storage support. A Mapping Strategy API is provided so that the user defines his/her own strategies. Nevertheless, we provide a set of predefined strategies to be used out-of-the-box. In our experiments we used four of these strategies to show how the selection of a strategy affects an algorithm's performance and cluster resources usage. Because the algorithm is a memory- and database-intensive algorithm with low CPU usage we selected two memory-based strategies, a location-based strategy and, as baseline for our experiments, a round-robin strategy. If the algorithm selected had an intensive use of CPU, for example, a content-based algorithm, a strategy that considers CPU characteristics such as number of cores and frequency, current and/or historical CPU usage, current and/or historical CPU load average, would have been more suitable. The strategies used in this work are listed below:

- *Available Memory* – this strategy uses jLiME's monitoring statistics to obtain the memory currently available on each node and then divides the given list of vertices accordingly. Clearly, this strategy is dynamic, that is, it adapts to the current status of the cluster.
- *Total Memory* – similarly to the Available Memory strategy, to divide the list, the Total Memory strategy uses the maximum amount of memory that a Peer can use. It is a fixed strategy that assigns more requests to nodes that have more memory available.
- *Location Aware* – this strategy takes advantage of the locality of the vertices. It divides the input into different lists of vertices grouped by their storage location. Such location is obtained applying a hashing function to each key. For example, let node *N1* be responsible for vertices *a1, a3, a5* and node *N2* for vertices *a2, a4*. If we use a Location Aware strategy to map vertices *a1, a2, a5*, it will divide the original list into: *a1, a5* → *N1* and *a2* → *N2*.
- *Round Robin* – this strategy simply divides the given list of vertices by the amount of nodes and assigns a sublist to each node. This equal division of requests among nodes makes this strategy the fairest strategy in terms of computing load. However, it does not consider either the locality of the data or the nodes characteristics such as available memory, CPU or physical network speed. This is a job mapping strategy commonly used in distributed computing as baseline for experiments.

As mentioned before, other strategies can be defined by the user. To support this feature, Graphly provides a Mapper interface that users can extend to define their own mapping strategies. The sole function of this interface is mapping an array of vertex IDs to the available cluster nodes. Additionally, a JobContext object is provided to access the cluster information and obtain shared objects from the platform. The RoundRobin Mapper implementation is shown in Figure 6 as an example.

Regarding the strategies that perform a division based on node information, the user may use the CriteriaMapper object. A CriteriaMapper divides the list of vertex IDs proportionally to a provided (quantitative) node metric such as available RAM memory, CPU speed, current amount of jobs being executed, among other metrics. For example, the Available Memory strategy uses a CriteriaMapper instantiated with a metric named 'memory.available'. Thus, using this strategy, the nodes that report larger amounts for the 'memory.available' metric will receive more vertex IDs than the other nodes. In our tests, we created our custom strategies using Spring Configuration.⁵ The Available Memory mapping strategy configuration is shown in Figure 7.

6. Case study

In order to validate the proposed architecture, we implemented a recommendation algorithm and performed experiments with different job distribution strategies in the current materialization of LMGPA. The algorithm used in this paper for illustrating the importance of mapping strategies is a followee recommender algorithm for Twitter users [18]. To simplify

```

public class RoundRobinMapper extends Mapper {
    @Override
    public Map map(int[] vertexList, JobContext env) {
        HashMap<ClientNode, TIntArrayList> map = new HashMap();
        ArrayList<ClientNode> serverList = env.getCluster().getExecutors();
        int cursor = 0;
        for (int vertex : vertexList) {
            ClientNode p = serverList.get(cursor);
            cursor = (cursor + 1) % serverList.size();
            map.get(p).add(vertex);
        }
        return map;
    }
}

```

Figure 6. Round Robin Mapping Strategy implementation.

```

<bean id="availablememory"
      class="edu.jlime.collections.adjacencygraph.mappers.CriteriaMapper">
    <constructor-arg>
        <bean class="edu.jlime.metrics.sysinfo.filter.SimpleInfoFilter">
            <constructor-arg value="sysinfo.mem.available" />
        </bean>
    </constructor-arg>
</bean>

```

Figure 7. Available Memory configuration using Spring Configuration.

the experimentation and analysis of strategy impact, we focus our experiments on a recommendation algorithm that only uses the topological information of the graph. However, the platform allows adding content to both vertices and edges, for example, ‘tweet lists’ and ‘type of relationship’, respectively. Naturally, algorithms that process content have memory and CPU requirements that are much higher than topological algorithms.

Section 6.1 presents the algorithm definition and its adaptation to the Graphly API. Section 6.2 describes the experimental settings, including the dataset used, the hardware features of the computer cluster and the initial configuration. In turn, Section 6.3 shows how the choice of a mapping strategy affects the algorithm execution times and cluster resource usage.

6.1. Follower recommendation algorithm

In the context of the Twitter social network, a user has a group of followees (outgoing edges) and followers (incoming edges). The algorithm used as case study, proposed in Armentano et al. [18], can be divided in two stages. It first explores the followee/follower network near the target user (i.e. the user receiving the suggestions) to select a set of potential users to recommend. Then, the candidates are filtered and ranked according to different criteria [56], such as the number of common friends with the target user or the matching of their content-based profiles. In this work we focus on supporting the first stage of the algorithm on a large graph, setting aside the quality of the actual recommendation.

The exploration of the followee/follower network in this algorithm is based on the characterization of Twitter users made in several studies [64, 65] and the fact that online social networks had become real-time information sources and news spreading mediums besides fostering the formation of social ties [66]. From an information-centric point of view, users are mainly divided into two categories: information sources and information seekers. Users behaving as information sources tend to collect a large number of followers as they post useful information or news, whereas information seekers follow several users to get information but rarely post a tweet themselves.

The rationale behind the graph traversal stage of this algorithm relies on the categorization of users as information seekers or information sources. Thus, it is assumed that the target user is an information seeker that has already

```

ListQuery followees = graph.get(user).ougoing();
                        VertexQuery   GetQuery

TopQuery query = followees.incoming().remove(followees)
                    GetQuery (removing original users from the result)

                    .countOutgoing().remove(followees).top(10);
                    CountQuery(counts users appearances and removes original users)   Filters the top 10 users from the count query.

Map<Integer,Integer> recommendation = query.query(cluster);
                                       Executes the whole query at a Server

```

Figure 8. The studied recommendation algorithm expressed using the Graphly API.

identified some interesting Twitter users acting as information sources (i.e. his/her followees). Other people who also follow these people (i.e. followers of the user followees) are likely to share some interests with the target user and might have discovered other relevant information sources on the same topics (i.e. their followees). This last group is then composed of potential candidates for suggesting to the user as future followees.

More formally, the search of candidate users is performed according to the following steps:

Step 1. Starting with the target user u_T (an information seeker), obtain the list of users he/she follows into a list $S = \{s_1, s_2, \dots, s_n\}$. Members of S are likely to be information sources.

Step 2. For each element in S , obtain its followers. Let us call the union of all these lists L , that is, $L = \cup_{s \in S} \text{followers}(s)$. Members of L are likely to be information seekers.

Step 3. For each element in L obtain its followees. Let us call the union of all these lists T , that is, $T = \cup_{l \in L} \text{followees}(l)$. Members of T are likely to be information sources.

Step 4. Exclude from T those users who the user being recommended is already following. Let us call the resulting list of candidates R .

Using the Graphly API, this algorithm can be expressed as depicted in Figure 8. The adapted version of the algorithm broadly contains the following steps:

- (1) Get the followees (*outgoing*) of the target user.
- (2) Get the followers (*incoming*) of the followees and remove the original followees from the list.
- (3) Count the followees (*countOutgoing*) of the followers and remove the original followees from the map.
- (4) Obtain the top n elements with most appearances ($n = 10$ in this example).

The *followees* object appears twice in the query because the original followees must not be part of the recommendation as they are already being followed by the user.

6.2. Experimental setting

Experiments were carried out using a Twitter dataset⁶ as of July 2009 consisting of approximately 1400 million relationships between more than 41 million users [66]. This is an extensively studied dataset [18, 67–71] containing topological

Table 1. Cluster hardware characteristics.

| | First set (five nodes) | Second set (three nodes) |
|------------------|--------------------------------|--------------------------|
| CPU | AMD Phenom II X6 1055T 2.8 GHz | AMD FX 6100 3.3 GHz |
| Number of cores | 6 | 6 |
| RAM | 8 GB | 16 GB |
| Physical network | 1 Gbit Ethernet | 1 Gbit Ethernet |
| Hard disk | 500 GB – 7200 rpm | 500 GB – 7200 rpm |

information about the social network, that is, it has binary relationships between user IDs, but it does not include tweets. This is not a limitation since the test algorithm builds recommendations based on user relationships only, searching for candidate followees in the Twitter graph. In addition, the dataset provides a good benchmark for topology-based link-prediction algorithms since it is a complete, real-world graph, not a subsample or artificially generated one. Thus, it is possible to tune the platform and observe its behaviour with real-world processing requirements.

For testing the algorithm, we selected two test groups of users out of the complete dataset with the goal of illustrating the performance of the different strategies for varying algorithm needs. For selecting a representative test user list, we first filtered the users using the information source ratio [19], denoting to what extent a user can be considered as an information source. This ratio is defined as follows:

$$IS(u) = \frac{followers(u)}{followers(u) + followees(u)} \quad (1)$$

We set this ratio to be ≤ 0.5 , which means that the number of followees had to be equal to or larger than the number of followers. This restriction arises as a natural way of selecting which users would have any interest in receiving recommendations, since the algorithm under study is designed for recommending followees to information consumers more than to information producers.

The first test group, called the ‘IS’ group, comprises the top 100 users with the smallest IS values. From the algorithm point of view, these users are the most interesting information consumers to recommend other users to follow. To obtain the second test group, called the ‘Followees’ group, we ordered the filtered users by number of followees and kept the top 100. These users represent the most influential users of the test set and were used to stress the computational nodes employed assuming that they will produce long lists of followers and followees.

Regarding the cluster characteristics, we used a heterogeneous cluster of eight nodes divided into two sets, each having identical hardware characteristics. Table 1 summarizes the most relevant characteristics of each set of nodes.

6.3. Results

The experiments were carried out as follows. For each user in the testing set, we ran the algorithm five times for each one of the strategies presented in Section 5.3. In every execution, we measured the total recommendation time, the bytes sent over the network and the maximum memory consumed (the biggest memory spike). Using this information we obtained the average number of bytes sent over the network, the average maximum memory consumption and the average recommendation time. It is worth noting that the number of times we ran the algorithms yielded as a result very low statistical deviations, and thus these averages are statistically significant. The distribution of average values for a given group of users and a given strategy is summarized in Figure 9.

As it was expected, users in the ‘IS’ group demand, on average, less time to process as well as less resources. Instead, the ‘Followees’ group was intentionally selected to stress the cluster. Users in this group need more resources to be processed and, consequently, their processing takes longer than users in the ‘IS’ group.

Regarding the performance of the mapping strategies for both groups of users, the Round Robin strategy provided the worst recommendation times as well as higher network and memory usage. This can be explained by the fair but random division of tasks processing.

The memory-based strategies provided similar performance in network and memory consumption and recommendation time. These strategies performed better than the Location Aware strategy and the Round Robin strategy. This might be explained by the memory-intensive nature of the algorithm used as case study: on each step, the algorithm generates exponentially larger subresults that are merged and allocated as input for the next step. Thus, an allocation of tasks according to memory capabilities of the nodes provides better performance by using less virtual memory from the

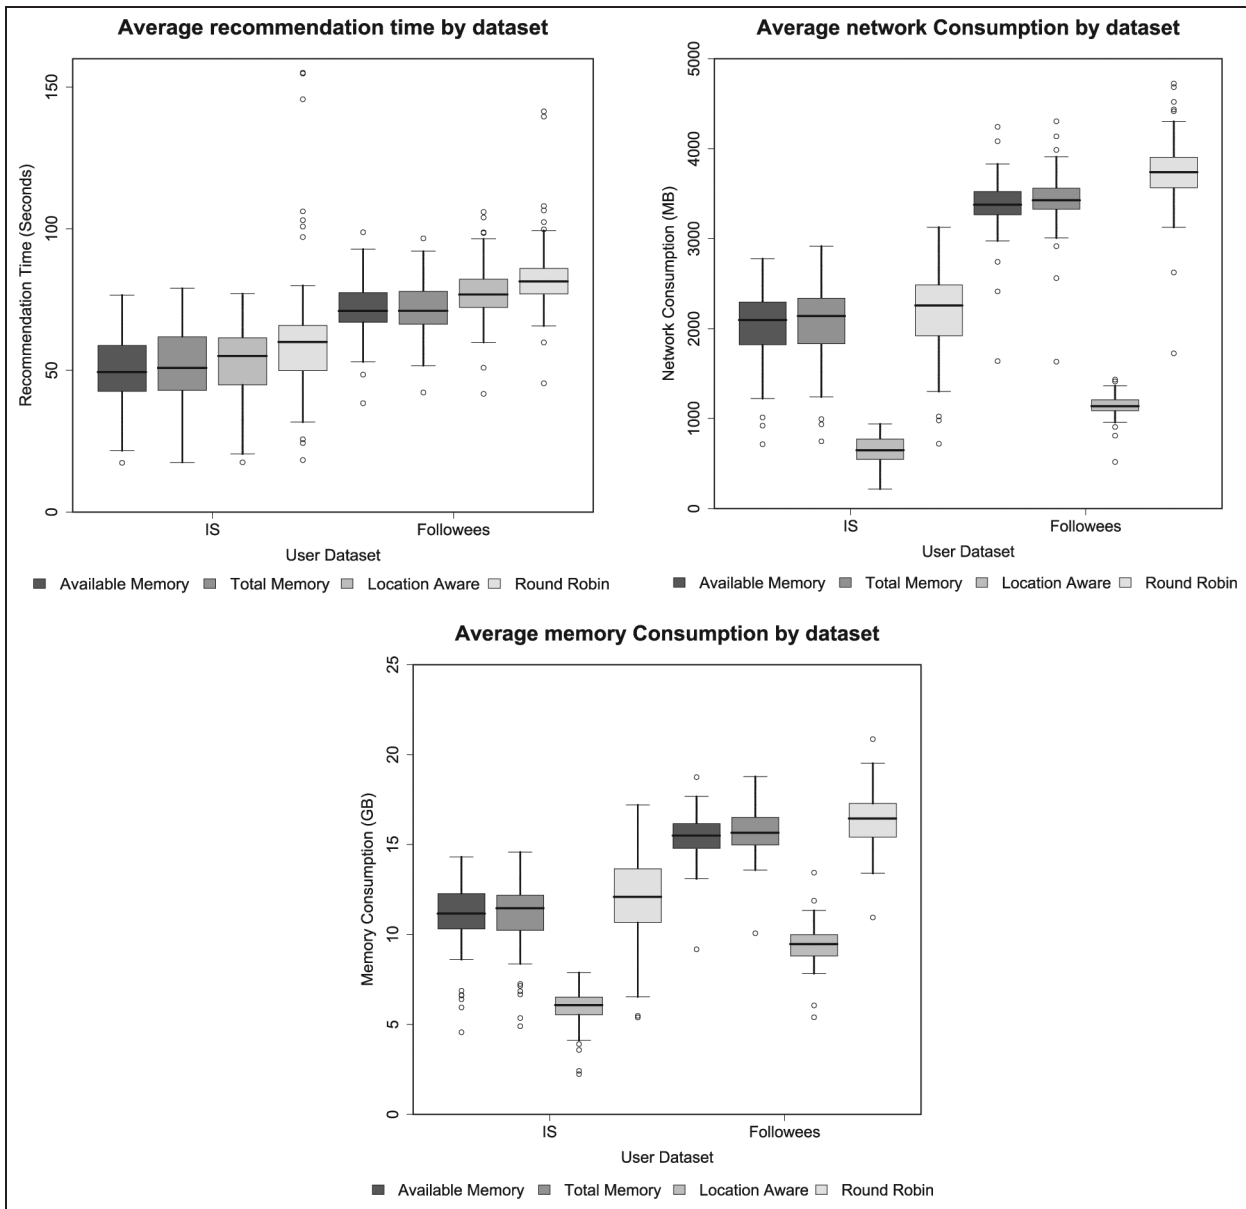


Figure 9. Average results for each testing group of users.

Operating System (also letting the OS to use more memory for disk cache) and, in the case of our implementation language, producing less garbage collection from the Java Virtual Machine.

Finally, the Location Aware strategy provided intermediate performance in terms of recommendation time. However, memory consumption was considerably lower than other strategies and, as expected, the network consumption was minimal. In scenarios where memory is scarce or network speed is an impediment, the Location Aware strategy may be a better fit than other strategies. Moreover, in some scenarios using Cloud processing platforms (e.g. Amazon AWS or Google Compute Engine), using less network bandwidth may have economic benefits.

To further compare strategies, a typical execution trace for the implemented recommendation algorithm is shown in Figure 10. In this case, the Location Aware strategy presents a memory usage pattern that is almost linear, whereas the other strategies show an exponential pattern on early steps of the recommendation algorithm. This is due to the nature of the algorithm, which generates an exponential amount of remote requests on each traversal step. As expected, the network usage is similar for non-locality strategies. Moreover, for those strategies, the growth in network usage matches a growth in memory usage.

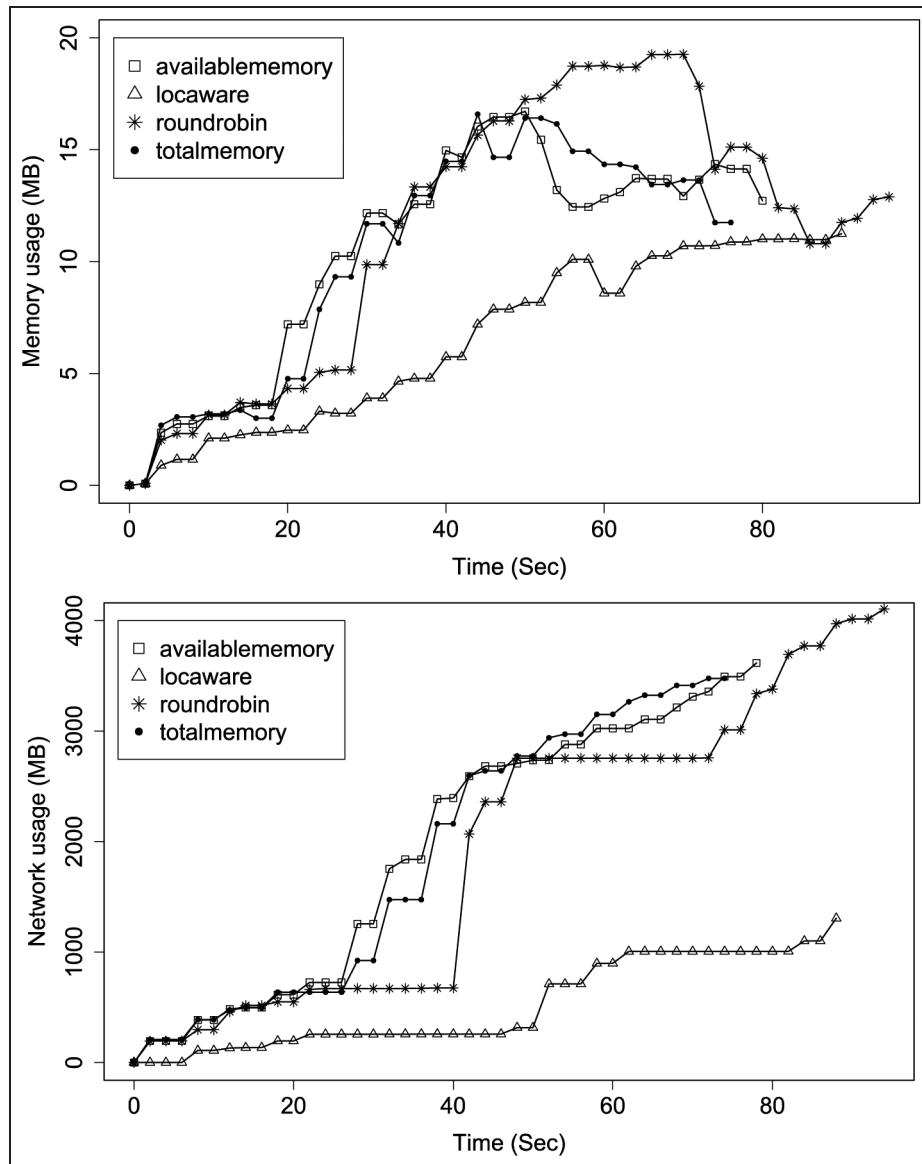


Figure 10. Cluster network and memory usage for an example user.

Although the platform allows developers of recommendation algorithms to decide the strategy according to the algorithm requirements, based on the previous results it is possible to draw some conclusions to guide this selection process. Table 2 provides assistance in the decision of which strategy shall be used depending on the environment characteristics, graph topology and algorithm characteristics. As environment characteristics we considered two scenarios: a Limited Memory scenario and a Slow/Costly Network scenario. The former corresponds to environments where the nodes have relatively low RAM memory, which is the case of, for instance, commodity clusters, or environments including embedded devices like cell phones or tablets [72] (a.k.a. mobile grids). The latter scenario considers slow computer networks, where transference between nodes incurs a big performance penalty or cases where network transfers are charged, as in, for example, paid Cloud Computing environments. Under graph topology we only consider graphs that are highly interconnected (e.g. natural graphs) and graphs that have a low amount of edges w.r.t. their vertices. Finally, in algorithm characteristics we identify two scenarios: implementing a memory-intensive recommendation algorithm or a CPU-intensive algorithm.

The scenarios depicted are far from exhaustive, mainly because we only focused on those scenarios that are related to the case study analysed. Then, even when the criteria used to fill in part of the table is somewhat subjective, our main

goal is to guide potential jLiME users into quickly obtaining an initial, working configuration for efficiently running their recommendation algorithms. Possible values for strategy performance in a given scenario range from very poor to very good. A very poor value does not necessarily mean that the current strategy is unusable for that scenario, but it might not perform as well as other strategies. Additionally, these values depend on the scenario being evaluated, for example, if we consider the Limited Memory scenario, a very poor value means that the strategy introduces high memory consumption, while a very good value means that the strategy favours low memory consumption.

The selection of a strategy for a given environment, graph and algorithm combination is straightforward. For example, in a Slow Network scenario the best choice is the Location Aware strategy. On the other hand, in a Limited Memory scenario, a low-connected graph and a CPU-intensive algorithm, the best choice might be a simple Round Robin strategy. Nevertheless, because the scenarios are not mutually exclusive, conflicts between requirements may arise, resulting in a trade-off situation.

7. Conclusions

In this paper, we presented a novel architecture and a corresponding implementation for designing distributed recommending algorithms. The algorithms are expressed in terms of graph traverse operations defined by the API of our graph support, Graphly. The underlying framework, jLiME, deals with networking concerns like remote code execution. The core motivation of this work is the creation of a customizable and transparent mechanism to distribute the processing of prototype recommending algorithms. This mechanism, called mapping strategies, works as a bridge between Graphly API and jLiME, by exposing information about the computational capabilities of every node, including the graph data distribution.

This work shows the importance of such mechanism by comparing several strategies and their impact in recommendation time and resource usage, especially on heterogeneous clusters. As a case study, we used an existing structural recommendation algorithm for Twitter and real Twitter data, and made recommendations to two different groups of 100 users to show the differences in recommendation times and resource usage. For this type of recommending algorithm, the memory-based strategies showed a very good recommendation time, followed by a strategy that uses data locality. This can, however, vary from algorithm to algorithm, and thus the architecture stressed the need for a flexible tuning support so that users can use or even define the strategies that best suit their algorithms.

As future work, we will examine the impact of mapping strategies in algorithms implemented using a different processing model, like Pregel. Some iterative graph algorithms, such as PageRank, can be better expressed in a vertex-centric processing model rather than a Divide-and-Conquer model like MapReduce or ForkJoin, that is, the ones that jLiME is currently based upon. Selecting a better model for a particular algorithm also results in performance benefits by reducing code and algorithmic overhead imposed by a less-fitted model. In the future, Table 2 (Section 6.3) might include more dimensions, such as the one representing the execution scheme (e.g. Pregel or Fork Join) that best suits algorithms w.r.t. how they traverse or process graphs. Under certain circumstances, a scheme would not suffice without a proper job mapping strategy. However, more experiments need to be conducted in the future to build such enhanced decision tables.

Furthermore, we will explore new mapping strategies and job-stealing strategies to rebalance processing if nodes are overloaded. While mapping strategies decide how task distribution is performed, job-stealing strategies perform a runtime adjustment on that distribution by stealing tasks that have not completed on busy nodes and allocating them on idle nodes. In other words, the current mapping strategies follow a push-based approach to distributing jobs to machines, whereas job stealing follows a pull-based approach. Indeed, these approaches have been studied in the context of other types of resource-demanding applications [73] and thus we believe they are worth exploring.

Another extension of this work includes the definition of reusable, higher-level graph primitive operations especially suited to recommending algorithms beyond simple graph primitives, like obtaining common neighbours, PageRank calculation or obtaining the ‘circle of trust’ – as described by the Who to Follow algorithm [6] – of a given vertex. The introduction of these new primitives might also require defining higher-level mapping strategies. A strategy that divides tasks according to the amount of neighbours of a given vertex might provide better load balancing when the primitive (or the user’s algorithm) is based on processing vertices neighbourhoods.

The inclusion of algorithms that use graph content besides topology will foster the creation of new strategies that take into account text mining requirements. For efficiency reasons, in many cases recommendation algorithms provide a preliminary ranking of users based on topological information, which is given as input to a content-based recommendation algorithm to build a final ranking. This approach offers a better balance of execution time–recommendation effectiveness, but still many job-mapping strategies could be explored at both ends to further speed up recommendations. Additionally, in terms of Graphly, these algorithms may become primitives that can be reused by other algorithms.

Table 2. Strategy suitability according to the environment, the graph and the algorithm characteristics.

| Strategies | Environment characteristics | | Graph topology | | Algorithm characteristics | |
|------------------|-----------------------------|---------------------|----------------|----------------------|---------------------------|---------------|
| | Limited memory | Slow/costly network | Natural graphs | Low-connected graphs | Memory-intensive | CPU-intensive |
| Available Memory | Poor | Very poor | Good | Good | Very good | Good |
| Total Memory | Poor | Very poor | Good | Good | Very good | Good |
| Location Aware | Very good | Very good | Good | Good | Good | Good |
| Round Robin | Very poor | Very poor | Poor | Good | Poor | Very good |

Funding

This work has been partially funded by ANPCyT, through project PICT-2011-0366, and CONICET, under grant PIP no. 114-200901-00381.

Notes

1. Giraph Web Page, <http://giraph.apache.org/>
2. jPregel Web Page, <http://kowshik.github.com/JPregel/>
3. GoldenOrb Web Page, <https://github.com/jzachr/goldenorb>
4. Phoebus Web Page, <https://github.com/xslogic/phoebus>
5. Spring Framework Web Page, <http://spring.io/>
6. <http://an.kaist.ac.kr/traces/WWW2010.html>

References

- [1] Durand G, Belacel N and LaPlante F. Graph theory based model for learning path recommendation. *Information Sciences* 2013; 251: 10–21.
- [2] Rausch K, Ntoutsi E, Stefanidis K and Kriegel H-P. Exploring subspace clustering for recommendations. In: *Proceedings of the 26th international conference on scientific and statistical database management (SSDBM '14)*, 2014, pp. 42:1–42:4.
- [3] Wang X, Ma J and Xu M. Group recommendation for Flickr images by 4-order tensor decomposition. *Journal of Computational Information Systems* 2014; 10(3):1315–1322.
- [4] Guo X and Lu J. Intelligent e-government services with personalized recommendation techniques. *International Journal of Intelligent Systems* 2007; 22(5): 401–417.
- [5] Jing Y, Zhang X, Wu L, Wang J, Feng Z and Wang D. Recommendation on Flickr by combining community user ratings and item importance. In: *Proceedings of the IEEE international conference on multimedia and expo (ICME 2014)*, 2014, pp. 1–6.
- [6] Gupta P, Goel A, Lin J, Sharma A, Wang D and Zadeh R. WTF: The who to follow service at Twitter. In: *Proceedings of the 22th International World Wide Web Conference*, 2013, pp. 505–514.
- [7] Lu J and Wang H. Variance reduction in large graph sampling. *Information Processing & Management* 2014; 50(3): 476–491.
- [8] Zikopoulos P and Eaton C. *Understanding big data: Analytics for enterprise class Hadoop and streaming data*. New York: McGraw-Hill Osborne Media, 2011.
- [9] Sakr S, Liu A, Batista DM and Alomari M. A survey of large scale data management approaches in cloud environments. *IEEE Communications Surveys Tutorials* 13(3): 311–336, 2011.
- [10] Malewicz G, Austern MH, Bik AJC, Dehnert JC, Horn I, Leiser N and Czajkowski G. Pregel: A system for large-scale graph processing. In: *Proceedings of the 2010 international conference on management of data (SIGMOD '10)*, 2010, pp. 135–146.
- [11] Krepeska E, Kielmann T, Fokkink W and Bal H. HipG: Parallel processing of large-scale graphs. *ACM SIGOPS Operating Systems Review* 2011; 45(2): 3–1.
- [12] Power R and Li J. Piccolo: Building fast, distributed programs with partitioned tables. In: *Proceedings of the 9th USENIX conference on operating systems design and implementation (OSDI '10)*, 2010, Vol. 10, pp. 1–14.
- [13] Low Y, Bickson D, Gonzalez J, Guestrin C, Kyrola A and Hellerstein JM. Distributed GraphLab: A framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment* 2012; 5(8): 716–727.
- [14] Dean J and Ghemawat S. MapReduce: Simplified data processing on large clusters. *Communications of the ACM* 2008; 51(1): 107–113.
- [15] Valiant LG. A bridging model for parallel computation. *Communications of the ACM* 1990; 33(8): 103–111.
- [16] Mateos C, Zunino A and Hirsch M. EasyFJP: Providing hybrid parallelism as a concern for divide and conquer Java applications. *Computer Science and Information Systems* 2013; 10(3): 1129–1163.

- [17] Mateos C, Zunino A and Campo M. JGRIM: An approach for easy gridification of applications. *Future Generation Computer Systems* 2008; 24(2): 99–118.
- [18] Armentano M, Godoy D and Amandi A. Topology-based recommendation of users in micro-blogging communities. *Journal of Computer Science and Technology* (Special Issue on Data Mining on Social Networks and Social Web) 2012; 27(3): 624–634.
- [19] Armentano M, Godoy D and Amandi A. Towards a followee recommender system for information seeking users in Twitter. In: *Proceedings of the international workshop on semantic adaptive social web (SASWeb'11)*, 2011.
- [20] Shao B, Wang H and Li Y. The Trinity Graph Engine, Microsoft Research, December, 2012, pp. 12–14, <http://research.microsoft.com/pubs/161291/trinity.pdf>.
- [21] Xin RS, Gonzalez JE, Franklin MJ and Stoica I. GraphX: A resilient distributed graph system on Spark. In: *Proceedings of the 1st international workshop on graph data management experiences and systems (GRADES '13)*, 2013, pp. 2:1–2:6.
- [22] Kleinberg JM. Authoritative sources in a hyperlinked environment. *Journal of the ACM* 1999; 46(5): 604–632.
- [23] Page L, Brin S, Motwani R and Winograd T. The PageRank citation ranking: Bringing order to the Web, Technical Report, Stanford InfoLab, 1999.
- [24] Lempel R and Moran S. SALSA: The stochastic approach for link-structure analysis. *ACM Transactions on Information Systems* 2001; 19(2): 131–160.
- [25] Morales-del-Castillo JM, Peis E, Ruiz AA and Herrera-Viedma E. Recommending biomedical resources: A fuzzy linguistic approach based on Semantic Web. *International Journal of Intelligent Systems* 2010; 25(12): 1143–1157.
- [26] Erling O and Mikhailov I. Virtuoso: RDF support in a native RDBMS. In: *Semantic web information management*, Berlin: Springer, 2010, pp. 501–519.
- [27] SYSTAP LLC. BigData, 2013, <http://www.blazegraph.com/bigdata> (accessed 14 April 2014).
- [28] Berners-Lee T, Hendler J and Lassila O. The Semantic Web. *Scientific American* 2001; 284(5): 28–37.
- [29] Harth A, Umbrich J, Hogan A and Decker S. YARS2: A federated repository for querying graph structured data from the Web. In: *The Semantic Web. Lectures Notes in Computer Science*, Vol. 4825. Berlin: Springer, 2007, pp. 211–224.
- [30] Think Aurelius. Titan. 2014, <http://thinkaurelius.github.io/titan/> (accessed 14 April 2014).
- [31] The Apache Foundation. HBASE, 2014, <http://hbase.apache.org/> (accessed 20 August 2014).
- [32] The Apache Foundation. Cassandra, 2014, <http://cassandra.apache.org/> (accessed 20 August 2014).
- [33] Oracle. Oracle BerkeleyDB, 2014, <http://www.oracle.com/technetwork/database/database-technologies/berkeleydb/overview/index.html> (accessed 20 August 2014).
- [34] Hazelcast Inc. Hazelcast, 2014, <http://hazelcast.com/> (accessed 20 August 2014).
- [35] Tinkerpop. Blueprints, 2014, <http://blueprints.tinkerpop.com/> (accessed 19 August 2014).
- [36] Orient Technologies. Orient DB, 2014, <http://orientdb.com/orientdb/> (accessed 20 August 2014).
- [37] Arango DB. Arango DB, 2014, <https://www.arangodb.com/> (accessed 20 August 2014).
- [38] Twitter Inc. FlockDB, 2013, <https://github.com/twitter/flockdb> (accessed 5 August 2013).
- [39] Topcuoglu H, Hariri S and Wu M. Performance-effective and low-complexity task scheduling for heterogeneous computing. *IEEE Transactions on Parallel and Distributed Systems* 2002; 13(3): 260–274.
- [40] Kim J-K, Sameer Shivle, Siegel HJ, Maciejewski AA, Braun TD, Schneider M, Tideman S, Chitta R, Dilmaghani RB, Joshi R, Kaul A, Sharma A, Sripada S, Vangari P and Yellampalli SS. Dynamic mapping in a heterogeneous environment with tasks having priorities and multiple deadlines. In: *Proceedings of the international parallel and distributed processing symposium (IPDPS'03)*, 2003.
- [41] Lu L and Zhou T. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications* 2011; 390(6): 1150–1170.
- [42] Li L, Qian L, Cheng J, Ma M and Chen X. Accurate similarity index based on the contributions of paths and end nodes for link prediction. *Journal of Information Science* 2015; 41(2): 167–177.
- [43] Li L, Ma M, Lei P, Leng M and Chen X. S2R&R2S: A framework for ranking vertex and computing vertex-pair similarity simultaneously. *Journal of Information Science* 2014; 40(6): 723–735.
- [44] Lü L, Jin C-H and Zhou T. Similarity index based on local paths for link prediction of complex networks. *Physics Review E* 2009; 80: 046122.
- [45] Schall D. Who to follow recommendation in large-scale online development communities. *Information and Software Technology* 2013; 56(12): 1543–1555.
- [46] Wu H, Sorathia V and Prasanna VK. Predict whom one will follow: Followee recommendation in microblogs. In: *Proceedings of the 2012 international conference on social informatics (SocialInformatics)*, 2012, pp. 260–264.
- [47] Zhang J, Wang Y and Vassileva J. SocConnect: A personalized social network aggregator and recommender. *Information Processing & Management* 2013; 49(3): 721–737.
- [48] Hannon J, McCarthy K and Smyth B. Finding useful users on twitter: Twittomender the followee recommender. In: *Advances in information retrieval. Lectures Notes in Computer Science*, Vol. 6611. Berlin: Springer, 2011, pp. 784–787.
- [49] Ying JJ-C, Lu EH-C and Tseng VS. Followee recommendation in asymmetrical location-based social networks. In: *Proceedings of the 2012 ACM conference on ubiquitous computing (UbiComp '12)*, 2012, pp. 988–995.
- [50] Chen T, Tang L, Liu Q, Yang D, Xie S, Cao X, Wu C, Yao E, Liu Z, Jiang Z, Chen C, Kong W and Yu Y. Combining factorization model and additive forest for collaborative followee recommendation. In: *KDD Cup Workshop*, 2012.

- [51] Li Y-M, Liao T-F and Lai C-Y. A social recommender mechanism for improving knowledge sharing in online forums. *Large-Scale and Distributed Systems for Information Retrieval. Information Processing & Management* 2012; 48(5): 978–994.
- [52] Buschmann F, Meunier R, Rohnert H, Sommerlad P, Stal M, Sommerlad P and Stal M. *Pattern-oriented software architecture: A system of patterns*. Chichester: John Wiley and Sons, 1996.
- [53] Tinkerpop. Gremlin, 2014 URL: <http://gremlin.tinkerpop.com/> (accessed 19 August 2014).
- [54] He H and Singh AK. Query language and access methods for graph databases. In: *Managing and mining graph data*, Vol. 40. Berlin: Springer, 2010.
- [55] Prud'Hommeaux E and Seaborne A. SPARQL query language for RDF, W3C recommendation, January 2008, Vol. 15.
- [56] Neo Technology Inc. Neo4J, 2013 (accessed 5 August 2013).
- [57] Varda K. Protocol buffers: Google's data interchange format, 2008, <http://google-opensource.blogspot.com/2008/07/protocol-buffers-googles-data.html> (accessed 19 August 2014).
- [58] The Apache Foundation. Apache Thrift, 2014, <https://thrift.apache.org/> (accessed 19 August 2014).
- [59] Downing TB. *Java RMI: Remote method invocation*. Chichester: John Wiley and Sons, 1998.
- [60] Videla A and Williams JJW. *RabbitMQ in action: distributed messaging for everyone*. Manning, 2012.
- [61] Ban B et al. JGroups, a toolkit for reliable multicast communication, 2002, <http://www.jgroups.org>
- [62] Google. LevelDB, 2014, <https://github.com/google/leveldb> (accessed 12 September 2014).
- [63] Corbellini A, Godoy D, Mateos C, Zunino A and Schiaffino S. A programming interface and platform support for developing recommendation algorithms on large-scale social networks. In: *Collaboration and technology*. Lectures Notes in Computer Science, Vol. 8658. Berlin: Springer, 2014, pp. 67–74.
- [64] Java A, Song X, Finin T and Tseng B. Why we Twitter: Understanding microblogging usage and communities. In: *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis (WebKDD/SNA-KDD '07)*, 2007, pp. 56–65.
- [65] Krishnamurthy B, Gill P and Arlitt M. A few chirps about Twitter. In: *Proceedings of the 1st workshop on online social networks (WOSP '08)*, 2008, pp. 19–24.
- [66] Kwak H, Lee C, Park H and Moon S. What is Twitter, a social network or a news media?. In: *Proceedings of the 19th international conference on World Wide Web (WWW'10)*, 2010, pp. 591–600.
- [67] Ediger D, Jiang K, Riedy J, Bader DA, Corley C, Farber R and Reynolds WN. Massive social network analysis: Mining Twitter for social good. In: *Proceedings of the 39th IEEE international conference on parallel processing (ICPP)*, 2010, pp. 583–593.
- [68] Lim KH and Datta A. Finding Twitter communities with common interests using following links of celebrities. In: *Proceedings of the 3rd international workshop on modelling social media, MSM '12*. New York: ACM, 2012, pp. 25–32, <http://doi.acm.org/10.1145/2310057.2310064>
- [69] Hu X, Tang L, Tang J and Liu H. Exploiting social relations for sentiment analysis in microblogging. In: *Proceedings of the sixth ACM international conference on Web search and data mining, WSDM '13*. New York: ACM, 2013, pp. 537–546, <http://doi.acm.org/10.1145/2433396.2433465>
- [70] Brown PE and Feng J. Measuring user influence on twitter using modified k-shell decomposition. In: *Proceedings of the fifth international AAAI conference on weblogs and social media*, 2011.
- [71] Yu SJ. The dynamic competitive recommendation algorithm in social network services. *Information Sciences* 2012; 187: 1–14, <http://www.sciencedirect.com/science/article/pii/S0020025511005718>
- [72] Rodriguez JM, Mateos C and Zunino A. Energy-efficient job stealing for CPU-intensive processing in mobile devices. *Computing* 2014; 96(2): 1–31.
- [73] Van Nieuwpoort RV, Wrzesińska G, Jacobs CJH and Bal HE. Satin: A high-level and efficient grid programming model. *ACM Transactions on Programming Languages and Systems* 2010; 32(3): 9:1–9:39.

Appendix A: jLiME Network Layer

Two default stack implementations, depicted in Figure A1, are provided by jLiME: an UDP-based stack and a TCP-based stack. Each stack might be used in different scenarios. TCP is a stream-oriented protocol, used when a reliable channel must be established between two networked processes.

On the other hand, UDP is packet or 'datagram'-oriented and does not provide any guarantee that after sending an array of bytes it will arrive at the destination. UDP has the advantage of consuming much less network resources than TCP as no connection needs to be established or maintained between nodes. In jLiME, however, we built a series of MPs to overcome UDP unreliability. First, to adjust to the UDP maximum datagram limit, we created a MessageFragmenter that divides the original packet into smaller pieces and assembles them on the receiving side. Second, an Acknowledge processor confirms received messages, and throttles sending if a given amount of confirmations are missing. This helps to maintain a steady stream of UDP packets by not overloading the receiver. Third, a Packet Bundler merges small messages in a bigger message or 'bundle' to reduce datagram overhead. Naturally, TCP already implements these techniques, but unlike the UDP stack, they are not optional.

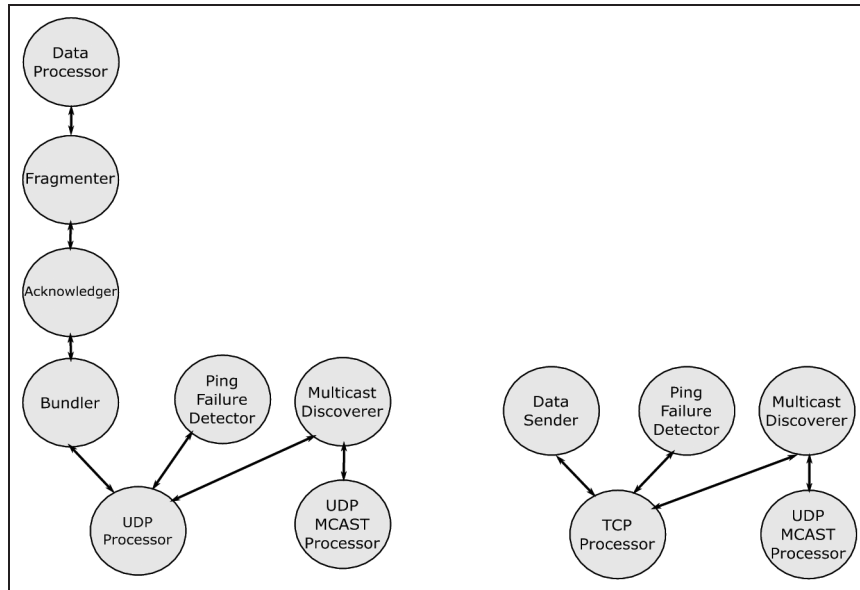


Figure 11. UDP and TCP stacks in jLiME.

Additionally, for both types of stack a custom discovery protocol was designed and implemented in two forms: Multicast Discovery and Ping Discovery. The Multicast variation sends UDP multicast messages to a given multicast address and expects responses from nodes. The Ping-based discovery simply sends UDP messages to a range of addresses and waits for responses. Similarly, a Ping-based protocol was created to track the status (alive or defunct) of nodes in the cluster. Lastly, streaming capabilities might be added to the jLiME stack by implementing the Streamer interface. The TCP-based implementation of this interface is straightforward as it is a stream-oriented protocol. The UDP Message Processor provides streaming capabilities by reassembling packets in the order they were sent.