Scalable Online Betweenness Centrality in Evolving Graphs

Nicolas Kourtellis, Gianmarco De Francisci Morales, Francesco Bonchi

Abstract—Betweenness centrality is a classic measure that quantifies the importance of a graph element (vertex or edge) according to the fraction of shortest paths passing through it. This measure is notoriously expensive to compute, and the best known algorithm runs in $\mathcal{O}(nm)$ time. The problems of efficiency and scalability are exacerbated in a dynamic setting, where the input is an evolving graph seen edge by edge, and the goal is to keep the betweenness centrality up to date. In this paper we propose the first truly scalable algorithm for online computation of betweenness centrality of both vertices and edges in an evolving graph where new edges are added and existing edges are removed. Our algorithm is carefully engineered with out-of-core techniques and tailored for modern parallel stream processing engines that run on clusters of shared-nothing commodity hardware. Hence, it is amenable to real-world deployment. We experiment on graphs that are two orders of magnitude larger than previous studies. Our method is able to keep the betweenness centrality measures up-to-date online, i.e., the time to update the measures is smaller than the inter-arrival time between two consecutive updates.

✦

1 INTRODUCTION

Betweenness centrality measures the importance of an element of a graph, either a vertex or an edge, by the fraction of shortest paths that pass through it [3, 14, 15].

Intuitively, an edge that connects two vertices that have many common neighbors is to some extent redundant. It belongs to a dense area of the graph, and information would be able to propagate even without it. In other terms, not many shortest paths will need such edge. This edge is what sociologists call a *strong tie*.

Conversely, a *weak tie* is an edge that connects two vertices with few common neighbors [16]. Such an edge is likely to bridge two distinct dense areas of the graph (also called *communities*), hence it participates in many shortest paths. Information has to go over this "bridge" to propagate from one community to another. The two vertices that own the bridge have a strategically favorable position because they can block information, or access it before the other individuals in their community: they span a *"structural hole"* [9].

Girvan and Newman [15] exploit this concept to define one of the first and most elegant algorithms for *community detection*. The algorithm iteratively removes the highest betweenness edge and produces a hierarchical decomposition of the graph, where the remaining disconnected components are the communities discovered.

Betweenness centrality has been used to analyze social networks [20, 26], protein networks [19], wireless ad-hoc networks [28], mobile phone call networks [10], multiplayer online gaming networks [2], to inform the design of socially-aware P2P systems [22], just to name a few examples.

• *N. Kourtellis, and F. Bonchi are with Yahoo Labs, Barcelona, Spain. G. De Francisci Morales is with Aalto University, Helsinki, Finland. E-mails: kourtell@yahoo-inc.com, gdfm@acm.org, bonchi@yahoo-inc.com*

Measuring betweenness centrality requires computing the shortest paths between all pairs of vertices in a graph. This computation is possible in small graphs with a few tens of thousands vertices and edges, but it quickly becomes prohibitively expensive as the graphs grow larger. Indeed, the best known algorithm for betweenness centrality, proposed by Brandes [6], runs in $\mathcal{O}(nm)$ time.

Due to its high cost, some works have proposed to parallelize its execution [4]. Others, have proposed to approximate betweenness centrality through the use of randomized algorithms [8, 23, 32]. However, the accuracy of these randomized algorithms can decrease considerably with the increase in graph size [23]. Variants of betweenness centrality, such as flow betweenness [13] and random-walk betweenness [30], also run in $\mathcal{O}(nm)$ time. Finally, there are no cheaper measures that can be used as a proxy, as they do not to correlate well with betweenness centrality [5] (differently from, e.g., degree centrality for PageRank). So there is no easy workaround to the complexity of computing betweenness centrality. State-of-the-art methods are too expensive for graphs with millions of vertices and edges, thus making this measure hard to use in practical application scenarios.

The picture gets even worse when considering that real graphs of interest, such as the Web, social networks, and information networks, are dynamic in nature and evolve continuously with new edges and vertices arriving, and old edges being removed. In such a scenario, the naïve approach of recomputing the measure from scratch is impractical even on moderately large graphs.

1.1 Related work

Recently there have been three main proposals by Lee et al. [25], Green et al. [17] and Kas et al. [21] for the incremental computation of vertex betweenness centrality.

QUBE [25] relies on the decomposition of the graph into disjoint minimum union cycles (MUCs). The algo-

rithm uses this decomposition to identify vertices whose centrality can potentially change. If the updates to the graph do not affect the decomposition, then the centrality needs to be recomputed (from scratch) only within the MUCs of the affected vertices. If the decomposition changes, a new one must be computed, and then new centralities must be computed for all the affected components. The performance of the algorithm is tightly connected to the size of the MUCs found, which in real (and especially social) graphs can be very large. The algorithm depends on a preprocessing step that computes a minimum cycle basis for the given graph. Then a MUC decomposition is computed by recursively taking the union of cycles in the minimum cycle basis that share a vertex. Therefore, it requires $\mathcal{O}(m)$ storage.

QUBE leverages Brandes' [6] algorithm to compute the centrality scores inside a MUC. However, the algorithm also requires to compute the (number of) shortest paths between each pair of vertices in the affected MUC, which could require $\mathcal{O}(n^2)$ space in the worst case. The overall space complexity is thus $\mathcal{O}(n^2 + m)$.

Green et al. [17] propose to maintain the previously computed values of betweenness and the data structures needed by Brandes' algorithm [6], and update the ones affected by a graph change. Their approach has a space complexity of $\mathcal{O}(n(n + m))$, which becomes prohibitive for large graphs of millions of vertices.

Kas et al. [21] extend the work proposed by Ramalingam and Reps [31] to accommodate the computation of vertex betweenness centrality while adding edges or vertices. Differently from Brandes' algorithm, their technique does not use dependencies but the actual shortest distances. It keeps the data structures needed to update the betweenness centrality of vertices: distances, number of shortest paths and predecessors list. The computational complexity can be at most as Brandes', i.e., $\mathcal{O}(nm)$. However, the space complexity is the same as Green et al. [17], i.e., $\mathcal{O}(n(n + m))$.

Recently, and concurrently with our study, some works have also improved the space complexity of the incremental betweenness computation to $O(n^2)$ [27, 29]. Similarly to our approach, McLaughlin et al. [27] modifies the technique by Green et al. [17] by dropping the predecessor lists, and porting it to GPUs to run on larger graphs. However, it has only been tested on a small sample of source nodes (256) instead of the full graph. Thus, its scalability to large graphs with higher memory demands, as well as its capability to follow dynamic rates of graph updates have not been demonstrated. Nasre et al. [29] present a tighter upper bound of $O(nm^*)$ on time complexity, where m^* is the number of edges that lie on shortest paths. This is the first algorithm to have a lower time complexity than Brandes' [6]. However, while usually smaller, m^* is still $O(m)$ in the worst case (e.g., an unweighted clique). Most importantly, their algorithm is not straightforward to parallelize, as it requires access to the shortest paths DAG (SPDAG) rooted in v from each SPDAG rooted in s ($\forall s \in V$), where v is an endpoint

TABLE 1: Comparison with previous studies: vertex (C_V) and edge centrality (C_E), edge addition (+) and removal (-), parallel and streaming computation ($\parallel$), size of the largest graph used in the experiments ($|V|$ and $|E|$). Note that Nasre et al. [29] have smaller time complexity than Brandes' and other algorithms.

Method	Year	Space	C_V	C_E	+	−	$\parallel$	$|V|$	$|E|$
Lee et al. [25]	2012	$\mathcal{O}(n^2+m)$	✓	✗	✓	✓	✗	12k	65k
Green et al. [17]	2012	$\mathcal{O}(n^2+nm)$	✓	✗	✓	✗	✗	23k	94k
Kas et al. [21]	2013	$\mathcal{O}(n^2+nm)$	✓	✗	✓	✗	✗	8k	19k
Nasre et al. [29]	2014	$\mathcal{O}(n^2)$	✓	✗	✓	✗	✗	-	-
This work	2014	$\mathcal{O}(n^2)$	✓	✓	✓	✓	✓	2.2M	5.7M

of the updated edge. Therefore, it is not possible to parallelize the algorithm by distributing the set of SPDAG on several machines, as in our approach. Moreover, neither of these approaches can handle removal of edges or provide updated edge betweenness scores.

1.2 Contributions

The main contribution of this paper is to provide *the first truly scalable and practical framework for computing vertex and edge betweenness centrality of large evolving graphs, incrementally and online*. Our proposal represents an advancement over the state of the art in four main aspects as summarized in Table 1.

First, our method maintains *both* vertex and edge betweenness centrality up-to-date for the same computational cost, while the previously proposed methods are only tailored for vertex betweenness.

Second, we handle *both additions and removals* of edges in a unified approach, while the previously proposed methods, besides QUBE, can handle only addition of edges. In fact, we show that the incremental, up-to-date edge betweenness under continuous edge removals allows for faster execution of the Girvan-Newman [15] algorithm on larger graphs.

Third, our method has reduced space overhead. Similarly to previous work [17], our algorithm maintains the previously computed values of betweenness and other needed data structures, and updates the ones affected by graph changes. However, our method avoids maintaining the predecessors lists, thus reducing the space complexity to $\mathcal{O}(n^2)$ [18]. This optimization requires a scan of all neighbors instead of predecessors: we show that this does not affect the time complexity, and makes the algorithm more scalable and faster in practice.

Forth, our framework is truly scalable and amenable to real-world deployment. The framework is carefully engineered to use *out-of-core* techniques to store its data structures on disk in a compact binary format. Data structures are read sequentially by employing *columnar storage*, and memory structures are *mapped directly on disk* to minimize memory copies.

Finally, we show how our method can be parallelized and deployed on top of modern parallel data processing engines that run on clusters of commodity hardware,

such as Storm[1], S4[2], Samza[3], or Hadoop. Our experiments test our method on graphs with millions of vertices and edges, i.e., two orders of magnitude larger than previous studies. By experimenting with real-world evolving graphs, we also show that our algorithm is able to keep the betweenness centrality measure up to date *online*, i.e., the time to update the measure is always smaller than the inter-arrival time between two consecutive updates.

An open-source implementation of our method is available on GitHub.[4] Inclusion in SAMOA,[5] a platform for mining big data streams [11, 12] is also planned.

2 PRELIMINARIES

Let $G = (V, E)$ be a (directed or undirected) graph, with $|V| = n$ and $|E| = m$. Let $P_s(t)$ denote the *set of predecessors* of a vertex t on shortest paths from s to t in G. Let $\sigma(s, t)$ denote the *total number of shortest paths* from s to t in G and, for any $v \in V$, let $\sigma(s, t \mid v)$ denote the number of shortest paths from s to t in G that *go through* v. Note that $\sigma(s, s) = 1$, and $\sigma(s, t \mid v) = 0$ if $v \in \{s, t\}$ or if v does not lie on any shortest path from s to t. Similarly, for any edge $e \in E$, let $\sigma(s, t \mid e)$ denote the number of shortest paths from s to t in G that *go through* e. The betweenness centrality of a vertex v is the sum over all pairs of vertices of the fractional count of shortest paths going through v.

Definition 2.1 (Vertex Betweenness Centrality): For every vertex $v \in V$ of a graph $G(V, E)$, its betweenness centrality $VBC(v)$ is defined as follows:

$$VBC(v) = \sum_{s,t \in V, s \neq t} \frac{\sigma(s, t \mid v)}{\sigma(s, t)}. \tag{1}$$

Definition 2.2 (Edge Betweenness Centrality): For every edge $e \in E$ of a graph $G(V, E)$, its betweenness centrality $EBC(e)$ is defined as follows:

$$EBC(e) = \sum_{s,t \in V, s \neq t} \frac{\sigma(s, t \mid e)}{\sigma(s, t)}. \tag{2}$$

Brandes' algorithm [6] leverages the notion of *dependency score* of a source vertex s on another vertex v, defined as $\delta_s(v) = \sum_{t \neq s, v} \frac{\sigma(s,t|v)}{\sigma(s,t)}$. The betweenness centrality $VBC(v)$ of any vertex v can be expressed in terms of dependency scores as $VBC(v) = \sum_{s \neq v} \delta_s(v)$. The following recursive relation on $\delta_s(v)$ is the key to Brandes' algorithm:

$$\delta_s(v) = \sum_{w:v \in P_s(w)} \frac{\sigma(s, v)}{\sigma(s, w)} (1 + \delta_s(w)) \tag{3}$$

The algorithm takes as input a graph $G=(V, E)$ and outputs the betweenness centrality $VBC(v)$ of every

---

Input: Graph $G(V, E)$ and edge update stream E_S
Output: $VBC'[V']$ and $EBC'[E']$ for updated $G'(V', E')$
   Step 1: Execute Brandes' alg. on G to create & store data structures for incremental betweenness.
   Step 2: **For each** update $e \in E_S$, execute Algorithm 1.
     Step 2.1 Update vertex and edge betweenness.
     Step 2.2 Update data structures in memory or disk for next edge addition or removal.

Fig. 1: The proposed algorithmic framework.

$v \in V$. It runs in two phases. During the first phase, it performs a search on the whole graph to discover shortest paths, starting from every source vertex s. When the search ends, it performs a *dependency accumulation* step by backtracking along the shortest paths discovered. During these two phases, the algorithm maintains four data structures for each vertex found on the way: a predecessors list $P_s[v]$, the distance $d_s[v]$ from the source, the number of shortest paths from the source $\sigma_s[v]$, and the dependency $\delta_s[v]$ accumulated when backtracking at the end of the search.

On unweighted graphs, Brandes' algorithm uses a breadth first search (BFS) to discover shortest paths, and its running time is $\mathcal{O}(nm)$. The space complexity of the algorithm is $\mathcal{O}(m + n)$. While this algorithm was initially defined only for vertex betweenness it can be easily modified to produce edge betweenness centrality at the same time [7].

3 FRAMEWORK OVERVIEW

Our framework computes betweenness centrality in evolving unweighted graphs. We assume new edges are added to the graph or existing edges are removed from the graph, and these changes are seen as a stream of updates, i.e., one by one. Henceforth, for sake of clarity, we assume an undirected graph. However, our framework can also work on directed graphs by following outlinks in the search phase and inlinks in the backtracking phase rather than generic neighbors.

The framework is composed of two basic steps shown in Figure 1. It accepts as input a graph $G(V, E)$ and a stream of edges E_S to be added/removed, and outputs, for an updated graph $G'(V', E')$, the new betweenness centrality of vertices (VBC') and edges (EBC') for each vertex $v \in V'$ and edge $e \in E'$.

The framework uses Brandes' algorithm as a building block in step 1: this is executed only once, offline, before any update. We modify the algorithm to (*i*) keep track of betweenness for vertices and edges at the same time, (*ii*) use additional data structures to allow for incremental computation, and (*iii*) remove the predecessors list to reduce the memory footprint and make out-of-core computation efficient.

Edge betweenness. By leveraging ideas from Brandes [7], we modify the algorithm to produce edge betweenness centrality scores. To compute simultaneously both edge and vertex betweenness, the algorithm stores the

intermediate dependency values (Eq. 3) independently for each vertex.

Additional data structures. To allow for incremental computation we need to maintain some additional data. In particular, we need a compact representation of the *directed acyclic graph of shortest paths rooted in the source vertex* (which we refer to as SPDAG), and the accumulated dependency values. Thus, for each source vertex s we maintain an additional data structure $BD[s]$ that stores its *betweenness data*. $BD[s]$ stores three pieces of information for each other vertex t:

- $BD[s].d[t]$: the distance of vertex t from source s;
- $BD[s].\sigma[t]$: the number of shortest paths starting from source s and ending at the given vertex t;
- $BD[s].\delta[t]$: the dependency accumulated on the vertex t in the backtracking to source s.

The data structure is initialized in step 1, and populated at the end of the dependency accumulation phase. Then, it is used in step 2.1 and updated in step 2.2.

Memory optimization. Brandes' algorithm builds a list of predecessors during the search phase to speed up the backtracking phase. Differently from the other data structures, the size of this list is variable and can grow considerably. For example, by assuming just 4 predecessors on average, the size of the list would be as large as $BD[\cdot]$ (assuming integer identifiers).

To reduce the space complexity of the algorithm, we remove the predecessors lists. When backtracking in the dependency accumulation phase, the algorithm checks all the neighbors of the current vertex, and uses the level of the vertex in the SPDAG (i.e., the distance from the source) to pick the next vertices to visit.

For each source, we need to maintain an SPDAG to all other vertices in the graph ($\mathcal{O}(n)$), along with the edges to their predecessors ($\mathcal{O}(m)$). Therefore, each SPDAG takes $\mathcal{O}(n+m)$ space. In total, the space complexity of the original algorithm is $\mathcal{O}(n(n+m))$ with the predecessors lists. By removing the predecessors lists, we reduce the space complexity to $\mathcal{O}(n^2)$. Furthermore, the time complexity remains unchanged, as shown next.

To understand why removing the predecessors lists does not increase the time complexity, consider that the order in which vertices are traversed is unchanged. Assume there are k edges in the predecessors lists overall. Brandes' algorithm checks each edge once during the search phase to populate the predecessors lists. In the backtracking phase only the edges in predecessors list are checked. Thus, overall the complexity of the algorithm is $O(m+k)$. In the worst case, k is of the same order as m, hence, the complexity of traversing all the predecessors lists is bounded by $\mathcal{O}(m)$. Scanning the neighbors of each vertex is also bounded by $\mathcal{O}(m)$. Thus, the algorithm's worst-case time complexity is unchanged.

An additional benefit of removing the predecessors lists is to also avoid the overhead of building it during the traversal. In practice, this optimization not only reduces the space complexity, but also decreases the

---

Algorithm 1: Incremental computation of vertex and edge betweenness when adding or removing an edge.

Input: $G(V, E)$, (u_1, u_2), $VBC[V]$, $EBC[E]$, $BD[V]$
Output: $VBC'[V']$, $EBC'[E']$, $BD'[V']$
Initialization: $EBC'[(u_1, u_2)] = 0.0$
Addition: $E' \leftarrow E \cup (u_1, u_2)$; **Deletion:** $E' \leftarrow E \setminus (u_1, u_2)$

```
1  for s ∈ V' do
2  |  u_L ← findLowest(u_1, u_2, BD[s])
3  |  u_H ← findHighest(u_1, u_2, BD[s])
4  |  dd ← BD[s].d[u_L] − BD[s].d[u_H]
5  |  if dd == 0 then
6  |  |  continue // same level addition/deletion
7  |  if dd ≥ 1 then
8  |  |  for r ∈ V do
9  |  |  |  σ'[r] = BD[s].σ[r]; d'[r] = BD[s].d[r]; δ'[r] = 0
10 |  |  |  t[r] ← N_T // not touched before
11 |  |  LQ[|V|] ← empty queues; Q_BFS ← empty queue
12 |  |  Q_BFS ← u_L
13 |  |  if addition then
14 |  |  |  if dd == 1 then
15 |  |  |  |  execute Alg. 2; // 0 level rise
16 |  |  |  if dd > 1 then
17 |  |  |  |  execute Alg. 4; // 1 or more level rise
18 |  |  else if deletion then
19 |  |  |  if u_L has predecessors then
20 |  |  |  |  execute Alg. 2; // 0 level drop
21 |  |  |  else execute Alg. 6 // 1 or more level drop
   |  Data Structures Update:
22 |  for r ∈ V' do
23 |  |  BD[s].σ[r] = σ'[r]; BD[s].d[r] = d'[r]
24 |  |  if t[r] ≠ N_T then BD[s].δ[r] = δ'[r]
```

---

average running time of the algorithm, as shown in Section 6 and in previous work [18, 1]. Moreover, by removing the predecessors lists, we do not need to maintain any variable-length data structure in our algorithm. This simplification allows us to use very efficient out-of-core techniques to manage $BD[\cdot]$ when its size outgrows the available main memory (see Section 5.1).

3.1 Addition and removal of edges

The addition of a new edge may cause structural changes in a graph, i.e., changes in the distance between vertices. Depending on the previous distance between the two newly-connected endpoints, these changes may bring some vertices closer to the current source s. Similarly, when an existing edge is removed structural changes may move the furthest endpoint away from the source.

Edge addition. In step 2 of the framework the Algorithm 1 is executed to update the betweenness centrality of the graph when a new edge (u_1, u_2) is added. The algorithm has access to the data structure $BD[\cdot]$ computed in step 1, and runs independently for each source. For a given source s, the algorithm uses $BD[s]$ to determine which endpoint of the new edge is closest to the current source s (denoted u_H) and which one is furthest (denoted u_L). Let $dd(u_L, u_H)$ denote the difference in distance from the source of the two endpoints before the addition, i.e., $dd(u_L, u_H) = d(s, u_L) - d(s, u_H)$ where d represents

shortest path distance. Since only one edge is added at a time, we simply denote this difference as dd.

Depending on how large dd is, a different type of update is needed. In particular, three cases can arise:

- $dd = 0$ (Proposition 3.1);
- $dd = 1$ (**0 level rise**, Section 4.1);
- $dd > 1$ (**1 or more levels rise**, Section 4.2).

The first case involves two vertices that are at the same distance from the source vertex.

Proposition 3.1: Given two vertices u_1 and u_2 such that they have the same distance from a source vertex s, and an edge $e = (u_1, u_2)$ that connects the two vertices, no shortest path from s to any other node in the graph passes trough the edge e, i.e., $d(s, u_1) = d(s, u_2) \implies \forall t \in V, \sigma(s, t \mid e) = 0$.

Proof: Omitted for brevity (see [17, 24]). □

No shortest path goes through the edge, no change occurs in the SPDAG, so the current source can be ignored.

In the second case, the new edge connects two vertices whose distance from the source differs only by one (Fig. 2a). Thus, this addition does not cause any structural change in the SPDAG, and all the distances remain the same. However, new shortest paths can be created due to the addition, and therefore the shortest paths and the dependencies of the graph must be updated.

In the third and most complex case, $dd > 1$, structural changes occur in the SPDAG (Fig. 2b depicts this case after the rise of u_L). In order to handle these changes properly, we introduce the concept of *pivot*.

Definition 3.2 (Pivot): Let s be the current source, let $d()$ and $d'()$ be the distance before and after an update, respectively, we define *pivot* a vertex $p_V \mid d(s, p_V) = d'(s, p_V) \land \exists w \in neighbors(p_V) \colon d(s, w) \neq d'(s, w)$.

Thus, a pivot is a vertex that, under an edge addition or removal, does not change its distance from the source s, but has neighbors that do so.

When $dd > 1$, we need to first compute the new distances by leveraging the pivots. Given that their distance has not changed, we can use them as starting points to correct the distances in the SPDAG. In the case of addition, all the pivots are situated in the sub-dag rooted in u_L, so we can combine the discovery of the pivot with the correction of the shortest paths. The different cases that can arise are discussed in detail in Section 4.2.

There exists also a fourth case: the new edge connects two previously disconnected components. This case degenerates into the case $dd = 1$. Indeed, no previous shortest path existed between the two disconnected components, so there is no structural change in the SPDAG.

Finally, new vertices arriving in the graph are handled simply by adding them to the source set V' with a zero VBC'. Then, for all sources, the new vertex is considered as u_L with $d[u_L] = d[u_H] + 1$, where u_H is the other endpoint of the incoming edge (therefore $dd = 1$).

Edge removal. In the case of an edge (u_1, u_2) removed from the graph, dd is at most one, as the two endpoints are connected before the removal. In this case, one of the
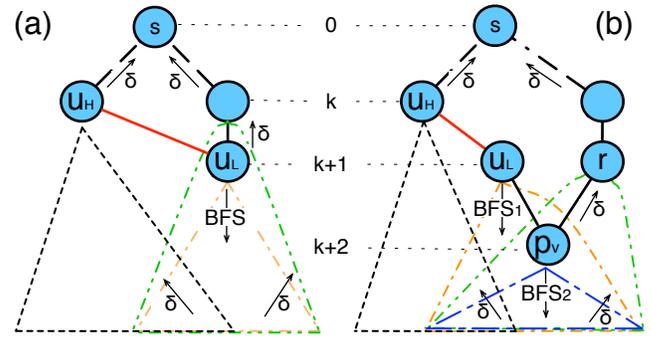

Fig. 2: The red (light) edge is added/removed and either does not cause structural changes (a), or does so (b).

two endpoints, u_H, is closest to the source, and clearly the edge (u_H, u_L) belongs to at least one shortest path from the source s to u_L. Therefore, the algorithm needs to check whether u_L has other shortest paths from s, not passing trough (u_H, u_L). Again, there are three cases:

- $dd = 0$ (Proposition 3.1);
- $dd = 1$ and u_L has other predecessors (**0 level drop**, Section 4.1);
- $dd = 1$ and u_L has no other predecessor (**1 or more levels drop**, Section 4.3).

In the second case, if u_L is connected to at least one vertex u'_H such that $dd(u_H, u'_H) = 0$, then u_L will remain at the same distance (Fig. 2a), and no structural change occurs. Thus distances remain the same. However, some shortest paths coming through (u_H, u_L) are lost, so the betweenness centrality needs to be updated.

In the third and most complex case, structural changes occur in the graph (Fig. 2b depicts this case before u_L drops). Also in this case we make use of pivots to correct the distances in the SPDAG first, and subsequently adjust the shortest paths and dependency values. However, not all pivots will be found in the sub-dag rooted in u_L after the removal. This difference makes this case more complicated than the addition, as some pivots cannot be discovered while adjusting the shortest paths (e.g., if nodes u_L and r were connected). Therefore, we need to first search and find the pivots, and then start a second BFS from those pivots to correct the shortest paths. The details of this case are covered in Section 4.3.

There is also the case where the edge removed disconnects the sub-dag rooted in u_L from the rest of the graph (or, similarly, turns u_L into a singleton). In this case, the shortest paths coming from the source, as well as the dependencies going to the source from this component must be removed and the betweenness adjusted. If u_L is to be removed, all its edges are iteratively removed and the singleton is replaced with zero VBC'.

4 INCREMENTAL ADDITION AND REMOVAL

In this section we discuss the details of our framework in the case of edge addition and removal.

Algorithm 2: Betweenness update for addition or removal of an edge where u_L remains at the same level after the update.

BFS Traversal from u_L:
1. $LQ[d'[u_L]] \leftarrow u_L$; $t[u_L] \leftarrow D_N$
2. **If** $addition$ **then** $\sigma'[u_L] += BD[s].\sigma[u_H]$;
3. **If** $deletion$ **then** $\sigma'[u_L] -= BD[s].\sigma[u_H]$;
4. **while** Q_{BFS} not $empty$ **do**
5. $v \leftarrow Q_{BFS}$
6.   **for** $w \in neighbors(v)$ **do**
7.   **if** $d'[w] == d'[v] + 1$ **then**
8.   **if** $t[w] == N_T$ **then**
9. $t[w] \leftarrow D_N$; $Q_{BFS} \leftarrow w$; $LQ[d'[w]] \leftarrow w$
10. $\sigma'[w] += \sigma'[v] - BD[s].\sigma[v]$

Dependency Accumulation:
11. **if** $deletion$ **then**
12. $\delta'[u_H] = BD[s].\delta[u_H] - \frac{\sigma[u_H]}{\sigma[u_L]}(1 + BD[s].\delta[u_L])$
13. $LQ[d'[u_H]] \leftarrow u_H$; $t[u_H] \leftarrow U_P$;
14. $level = |V'|$
15. **while** $level > 0$ **do**
16.   **while** $LQ[level]$ not $empty$ **do**
17. $w \leftarrow LQ[level]$
18.   **for** $v \in neighbors(w)$ **do**
19.   **if** $d'[v] < d'[w]$ **then**
20.   Execute module in Alg. 3.
21.   **if** $addition$ **then**
22.   **if** $t[v] == U_P$ **and** $(v \neq u_H$ **or** $w \neq u_L)$ **then**
23. $\delta'[v] -= \alpha$
24.   **if** $(v,w) \neq (u_L, u_H)$ **then** $EBC'[(v,w)] -= \alpha$
25.   **if** $deletion$ **then**
26. $EBC'[(v,w)] -= \alpha$
27.   **if** $t[v] == U_P$ **then** $\delta'[v] -= \alpha$
28.   **if** $w \neq s$ **then** $VBC[w] += \delta'[w] - BD[s].\delta[w]$
29. $level = level - 1$;

Algorithm 3: Initialization dependency module.

---

1. **if** $t[v] == N_T$ **then**
2. $t[v] \leftarrow U_P$; $\delta'[v] = BD[s].\delta[v]$; $LQ[level-1] \leftarrow v$
3. $c = \frac{\sigma'[v]}{\sigma'[w]}(1 + \delta'[w])$; $\delta'[v] += c$; $EBC'[(v,w)] += c$
4. $\alpha = \frac{BD[s].\sigma[v]}{BD[s].\sigma[w]}(1 + BD[s].\delta[w])$

4.1 No level change

Algorithm 2 handles the case when the edge added or removed does not cause structural changes in the SPDAG, i.e., $d(s, u_L)$ is unchanged (Fig. 2a). The algorithm initializes u_L by adding or removing all shortest paths from its predecessor u_H, depending on whether the edge (u_H, u_L) is added or removed. It also maintains a state flag $t[v]$ for each vertex v which represents the direction in which the algorithm has encountered the vertex v: D_N if descending (searching), U_P if ascending (backtracking), and N_T if untouched.

The vertices whose shortest paths from a source s have potentially been altered are situated in the sub-dag rooted in u_L. Thus, Algorithm 2 performs a BFS traversal of the SPDAG starting from u_L by visiting neighbors whose distance from s is higher than the current vertex (line 7, orange single-dotted triangle in Fig. 2a). During

the BFS, it updates the number of shortest paths to each vertex found to take into account the paths created (lost) by the addition (removal) of the edge (u_H, u_L). To avoid double counting, the old number of shortest paths from each predecessor $BD[s].\sigma[v]$ is subtracted before adding the new number of shortest paths $\sigma'[v]$ (line 10). The vertices encountered in the BFS are added to a level-specific queue LQ for later use. At the end of the BFS, the number of shortest paths is up-to-date.

In the dependency accumulation phase, the algorithm polls the LQ queue for each level and visits the vertices in reverse order of discovery by the BFS (δ arrows in Fig. 2a). In the case of edge removal (lines 11-13), u_H is inserted in LQ before the dependency accumulation starts, to guarantee that the data structures of u_H and its predecessors will be updated even though the edge (u_H, u_L) does not exist anymore.

The vertices whose dependency needs an update are situated in two parts of the SPDAG. Part of them is in the sub-dag rooted in u_L, as discovered by the BFS. For these vertices, the number of shortest paths was updated during the BFS. They have been inserted in the appropriate queue LQ and will be examined during the dependency accumulation phase. The others are the predecessors of the sub-dag, which lay at the fringe of the first ones (green, double-dotted triangle in Fig. 2a). These vertices are found while backtracking in the dependency accumulation phase by examining the neighbors of vertices in LQ. Discovered predecessors are added in LQ and examined in the next level (line 2 in Alg. 3).

During the backtracking phase, the algorithm corrects the dependency of vertices that are affected. If a vertex was not encountered before, the old dependency for the particular path is subtracted (line 23 for addition, line 27 for removal). The new dependency c is added in line 3 of Alg. 3 and then corrected by subtracting the old dependency α in line 24 for addition and line 26 for removal. Similarly the VBC is updated in line 28.

4.2 Addition: 1 or more levels rise

In this case, the SPDAG undergoes structural changes. To handle these changes, we need to find the pivots, which define the boundary of the structural change of the SPDAG. Given that the addition of the edge can only bring vertices closer to the source if they are reachable from u_L, all the pivots can be found while traversing the SPDAG with a BFS starting from u_L, similarly to the previous case (orange, single-dotted triangle in Fig. 2b).

Algorithm 4 shows the details. The algorithm begins by initializing the new distance of u_L (line 1). Similarly to the previous case, it adds all the new shortest paths to vertices found during its BFS traversal. However, in this case there are structural changes in the SPDAG due to the connection of two endpoints previously far apart. The vertices reachable from u_L might be pulled closer to the source s along with u_L. As a result, new shortest paths may emerge, old shortest paths may become obsolete,

Algorithm 4: Betweenness update for addition of an edge where u_L rises one or more levels after addition.

BFS Traversal from u_L:
1   $LQ[d[u_L]] \leftarrow u_L$; $d'[u_L] = BD[s].d[u_H] + 1$
2   **while** Q_{BFS} *not empty* **do**
3     $v \leftarrow Q_{BFS}$; $t[v] \leftarrow D_N$; $\sigma'[v] = 0$
4   **for** $w \in neighbors(v)$ **do**
5   **if** $d'[w] + 1 == d'[v]$ **then** $\sigma'[v] += \sigma'[w]$;
6   **if** $d'[w] > d'[v]$ **and** $t[w] == N_T$ **then**
7         $t[w] \leftarrow D_N$; $d'[w] = d'[v] + 1$; $LQ[d'[w]] \leftarrow w$; $Q_{BFS} \leftarrow w$
8   **if** $d'[w] == d'[v]$ **and** $BD[s].d[w] \neq BD[s].d[v]$ **then**
9   **if** $t[w] == N_T$ **then**
10           $t[w] \leftarrow D_N$; $LQ[d'[w]] \leftarrow w$; $Q_{BFS} \leftarrow w$
Dependency Accumulation:
11   $level = |V'|$; $te[e] \leftarrow N_T, e \in E$
12   **while** $level > 0$ **do**
13   **while** $LQ[level]$ *not empty* **do**
14       $w \leftarrow LQ[level]$
15   **for** $v \in neighbors(w)$ **do**
16   **if** $d'[v] < d'[w]$ **then**
17   Execute module in Alg. 3.
18   **if** $(t[v] = U_P)$ **and** $(v \neq u_H$ **or** $w \neq u_L)$ **then**
19             $\delta'[v] -= \alpha$
20   **if** $BD[s].d[v] == BD[s].d[w]$ **then** $\alpha = 0.0$
21   **if** $BD[s].d[w] < BD[s].d[v]$ **then**
22             $\alpha = \frac{BD[s].\sigma[w]}{BD[s].\sigma[v]}(1 + BD[s].\delta[v])$
23   **if** $(v, w) \neq (u_L, u_H)$ **then** $EBC'[(v, w)] -= \alpha$
24   **if** $d'[v] == d'[w]$ **and** $BD[s].d[w] \neq BD[s].d[v]$ **then**
25   Execute module in Alg. 5.
26   **if** $w \neq s$ **then** $VBC[w] += \delta'[w] - BD[s].\delta[w]$;
27     $level = level - 1$;

Algorithm 5: EBC correction if endpoints were not at the same level before the change.

1   **if** $te[(v, w)] == N_T$ **then**
2     $te[(v, w)] \leftarrow U_P$; $\alpha = 0$
3   **if** $BD[s].d[w] > BD[s].d[v]$ **then**
4       $\alpha = \frac{BD[s].\sigma[v]}{BD[s].\sigma[w]}(1 + BD[s].\delta[w])$
5   **if** $BD[s].d[w] < BD[s].d[v]$ **then**
6       $\alpha = \frac{BD[s].\sigma[w]}{BD[s].\sigma[v]}(1 + BD[s].\delta[v])$
7     $EBC[(v, w)] -= \alpha$

---

and distances from the source may change. Therefore, the vertices do not inherit the shortest paths from their predecessors (line 3), rather, the shortest paths are computed during the modified BFS.

The structural changes that can happen in the SPDAG are depicted in Figure 3. Let us examine these cases for a vertex x and its neighbor y. Let a *sibling* be a neighbor of vertex that is at the same distance from the source. Before the addition, x and y could be either siblings (case 1, Fig. 3) or predecessor and successor (case 2). If y is now a predecessor of x (case 1a), the algorithm adjusts the shortest paths of x (line 5). If x was and still is a predecessor of y (line 6), the new edge has caused both x and y to move closer to s by the same amount (case 2a). In this case, we update the distance from s and insert y in the BFS queue for further exploration (line 7). If y is now
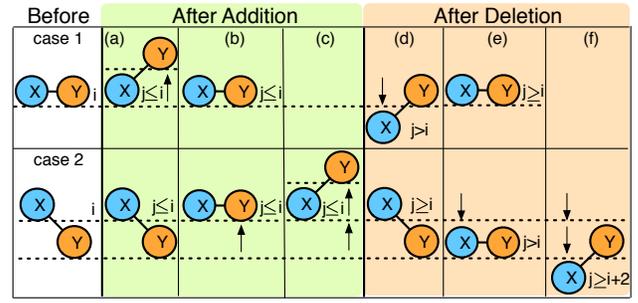

Fig. 3: Possible configurations of an edge before and after an update that causes structural changes.

on the same level as x, but was not before the addition (case 2b), y is added to the BFS for further exploration (line 10). If y moved two levels w.r.t. x (case 2c), it will be discovered first in the BFS after the update (line 6).

Clearly there are no structural changes in the vertices at levels above u_L (i.e., closer to the source). The possible sub-cases examined cover all possible scenarios of how a pair of connected vertices (and thus their edge) can be found after the addition of the new edge. The shortest paths and distances ($\sigma'[\cdot]$, $d'[\cdot]$) are updated in the way that the original Brandes' algorithm proposes.

In the dependency accumulation phase, the dependency score of all vertices examined is updated with the new number of shortest paths computed in the BFS phase. This part of the algorithm is similar to the corresponding one in Algorithm 2. However, there are important differences in the correction of the dependency for the edge betweenness centrality (lines 20–25, Alg. 4). Assuming v is x and w is y, if both x and y remain at the same relative distance from the source, the dependency to be subtracted α is calculated in line 4 of Alg. 3 (case 2a). However, if y moves closer (case 2c), then y was a successor of x but now it is a predecessor of x. Therefore, we need to subtract the dependency on y. The subtracted value is adjusted by switching w with v in the dependency accumulation formula (lines 21–22, Alg. 4).

If the endpoints of the edge were at the same level before the addition (case 1) there is no need for correction since no dependency was accumulated on the edge (line 20, Alg. 4). If the endpoints are now at the same level but were not before (case 2b), the old dependency needs to be subtracted from the betweenness of the edge. Also, the edge is marked not to be traversed again (Alg. 5). In Alg. 5, if w was a successor of v, the old dependency is calculated on line 4, whereas if w was a predecessor of v, the old dependency is calculated on line 6. The vertex betweenness centrality is updated on line 26 of Alg. 4 by adding the new dependency accumulated on the vertex w and subtracting the old dependency.

In summary, all possible cases of structural changes in the SPDAG below u_L are covered by Alg. 4, which correctly updates the betweenness scores and accompanying data structures of all affected vertices and edges.

Algorithm 6: Pivot finder when an edge removal causes u_L to drop one or more levels.

1   $PQ[|V|] \leftarrow$ empty queues; $Q_{BFS} \leftarrow$ empty queue
2   $first = V$; $Q_{BFS} \leftarrow u_L$; $t[u_L] \leftarrow N_P$;
3   **while** Q_{BFS} *is not empty* **do**
4     $v \leftarrow Q_{BFS}$
5   **for** $w \in neighbors(v)$ **do**
6   **if** $BD[s].d[w]+1 == BD[s].d[v]$ **and** $t[w] == N_T$ **and** $t[v] \neq P_V$ **then**
7         $PQ[d[v]] \leftarrow v$; $t[v] \leftarrow P_V$ // a new pivot
8   **if** $first > BD[s].d[v]$ **then**
9           $first = BD[s].d[v]$ // the first pivot
10   **else if** $BD[s].d[w] == BD[s].d[v]+1$ **or** $BD[s].d[w] == BD[s].d[v]$ **then**
11   **if** $t[w] == N_T$ **then**
12           $Q_{BFS} \leftarrow w$; $t[w] = N_P$ // not a pivot
13   **if** $PQ[.]$ *not empty* **then**
14   Execute Algorithm 7 // ≥ 1 pivots
15   **else if** $PQ[.]$ *is empty* **then**
16   Adjust appropriately for disconnected component.

4.3 Removal: 1 or more levels drop

Algorithm 6 shows how to handle an edge removal that causes u_L to move one or more hops away from the source s. First, the algorithm needs to find the pivots in the sub-dag rooted in u_L. Then, distances from the source are updated by exploring the graph starting from the discovered pivots. Finally, the usual dependency accumulation corrects the values of dependency and betweenness centrality.

The algorithm starts a BFS from u_L to search for vertices that have predecessors still connected to the rest of the graph, i.e., outside the scope of the sub-dag rooted in u_L (lines 6–9) and marks them as pivots (P_V) (BFS_1, orange, single-dotted triangle in Fig. 2b). The rest are non-pivots (N_P). Figure 3 illustrates the possible cases that can be encountered by this BFS traversal in the case of edge removal. Pivots are represented by vertex y if y remained in level i (case 1d), or if y remained in level $i+1$ (cases 2e and 2f).

After the first BFS finishes, Algorithm 7 starts a new BFS (BFS_2, blue, single-dashed triangle in Fig. 2b) from the pivots found. This new BFS first corrects the distances of the vertices discovered (line 9), and their shortest path counts (line 12). Furthermore, in lines 23-31 it adjusts the dependency and betweenness values in a similar fashion as described in Section 4.2, by covering all possible cases (Fig. 3) and following the δ arrows in Fig. 2b.

If u_L has at least one sibling before the removal of the edge, we can use the same general technique presented above. However, some optimizations to reduce the computation overhead are possible. Indeed, the first search from u_L to find the pivots can be avoided. In fact, in this case all possible scenarios shown in Figure 3 can be seamlessly found and resolved while adjusting the shortest paths (Algorithm 7), since the starting pivots are the siblings of u_L, and the other pivots are all found

Algorithm 7: Betweenness update for edge removal where u_L drops one or more levels after removal.

  BFS Traversal: from first pivot point(s)
1   $Q_{BFS} \leftarrow$ empty queue; $Q_{BFS} \leftarrow PQ[first]$
2   $next = first + 1$
3   **while** Q_{BFS} *not empty* **do**
4     $v \leftarrow Q_{BFS}$; $t[v] \leftarrow D_N$; $LQ[d'[v]] \leftarrow v$; $\sigma'[v] = 0$
5   **if** $next == d'[v] + 1$ **then**
6       $Q_{BFS} \leftarrow PQ[next]$; $next = next + 1$
7   **for** $w \in neighbors(v)$ **do**
8   **if** $t[w] == N_P$ **then**
9         $t[w] = D_N$; $d'[w] = d'[v] + 1$; $Q_{BFS} \leftarrow w$
10   **else if** $t[w] == P_V$ **then** $t[w] \leftarrow D_N$
11   **else**
12   **if** $d'[w] + 1 == d'[v]$ **then** $\sigma'[v] += \sigma'[v]$
13   **if** $d'[w] == d'[v]$ **and** $BD[s].d[w] \neq BD[s].d[v]$ **then**
14   **if** $BD.d[w] > BD.d[v]$ **and** $t[w] \neq D_N$ **then**
15             $t[w] \leftarrow D_N$; $LQ[d'[w]] \leftarrow w$; $Q_{BFS} \leftarrow w$
  Dependency Accumulation:
16   $\delta'[u_H] = BD[s].\delta[u_H] - \frac{\sigma[u_H]}{\sigma[u_L]}(1 + BD[s].\delta[u_L])$
17   $LQ[d'[u_H]] \leftarrow u_H$; $te[e] \leftarrow N_T, e \in E$; $t[u_H] \leftarrow U_P$; $level = |V'|$
18   **while** $level > 0$ **do**
19   **while** $LQ[level]$ *not empty* **do**
20       $w \leftarrow LQ[level]$
21   **for** $v \in neighbors(w)$ **do**
22   **if** $d'[v] < d'[w]$ **then**
23   Execute module in Alg. 3; $\alpha = 0$;
24   **if** $BD[s].d[w] > BD[s].d[v]$ **then**
25             $\alpha = \frac{BD[s].\sigma[v]}{BD[s].\sigma[w]}(1 + BD[s].\delta[w])$
26   **else if** $BD[s].d[w] < BD[s].d[v]$ **then**
27             $\alpha = \frac{BD[s].\sigma[w]}{BD[s].\sigma[v]}(1 + BD[s].\delta[v])$
28   **if** $t[v] == U_P$ **then** $\delta'[v] -= \alpha$
29           $EBC'[(v,w)] -= \alpha$
30   **if** $d'[v] == d'[w]$ **and** $BD[s].d[w] \neq BD[s].d[v]$ **then**
31   Execute module in Alg. 5.
32   **if** $w \neq s$ **then** $VBC[w] += \delta'[w] - BD[s].\delta[w]$
33     $level = level - 1$

during the BFS. These optimizations are explained in detail in an extended version of our work [24].

5 SCALABILITY

The algorithm described previously is able to update the betweenness centrality of a graph incrementally. Nevertheless, to attain our goal of a useful and practical framework for real-world deployment, the algorithm is not enough. Indeed, a number of scalability issues needs to be addressed in order to have a practical tool.

The space complexity of the algorithm is quadratic in the number of vertices. When dealing with large graphs, the space requirements can easily outgrow the available main memory. In this case, we can use the disk to store the required data structures, as our algorithm allows for an efficient out-of-core implementation.

Despite this feature, the space requirements for very large graphs can still outgrow the disk. Furthermore, disk access will become a bottleneck due to reading and writing large amounts of data. A simple solution is to

divide the execution across multiple machines with multiple disks. This solution not only allows the framework to scale to larger inputs, but also leads to improved speedup over the sequential version.

5.1 Out-of-core organization

The removal of the predecessors lists from the algorithm reduces its space complexity by $\mathcal{O}(nm)$. As an additional benefit, it leaves the algorithm with no variable-length data structure. Indeed, our algorithm stores only three fixed size data structures per vertex: the distance from the source $d[\cdot]$ (1 byte), the number of shortest paths from the source $\sigma[\cdot]$ (2 bytes) and the accumulated dependency $\delta[\cdot]$ (8 bytes). This organization allows an efficient data layout on disk. The graph is loaded and kept in memory, allowing fast random access of vertex neighborhoods.

We encode $BD[\cdot]$ in binary format on disk. For each source s, we store the data for each other vertex in a columnar fashion, i.e., we store on disk all the distances, then all the numbers of shortest paths, and finally the dependency values, in order: $\{BD[s].d[\cdot], BD[s].\sigma[\cdot], BD[s].\delta[\cdot]\}$. We avoid storing the vertex IDs for sources and destinations by storing the data structures sequentially on disk, and inferring the ID from the order. Overall, for each source the algorithm stores 3 arrays, with sizes $1 \times n$, $2 \times n$ and $8 \times n$ bytes, respectively. Because of the binary format, each array can be read by using file channels and byte buffers, loaded directly in memory and be ready for use. This optimization avoids memory allocations and memory copies during the execution of the algorithm. In practice, the computation happens at the speed of sequential disk access. In the future we plan to explore compression schemes to reduce the space and disk access overhead.

As explained in the previous sections, the work done by the algorithm depends on $dd(u_L, u_H)$. Therefore, after loading the distances from disk, we check the distance for the endpoints u_H and u_L. If they are at the same distance ($dd = 0$), we skip directly to the next source without loading the rest of the data structure. This operation is efficient because the data structures have fixed size so the offset to skip to reach the beginning of the next source is constant. Otherwise, the arrays are loaded in memory and Algorithm 1 is executed. When a source is covered, the algorithm writes the arrays back to disk, in place and sequentially.

5.2 Parallelization

The out-of-core version presented in the previous section is slower than the in-memory one, but enables the framework to scale to large graphs. However, the space requirements for real graphs can be staggering (for a graph of 1M vertices we need $\approx$ 11TB of space). Although nowadays is fairly easy to have such an amount of disk storage available on servers, reading and writing this amount of data would incur a significant overhead.

To solve this issue, we take advantage of the parallel nature of our framework and propose to distribute the computation on a cluster of shared-nothing machines. Let p be the number of available machines in the system. We distribute the data structure $BD[\cdot]$ evenly among the p machines, i.e., each machine will be allocated $\approx n/p$ sources. This parallelization is possible since each source can be examined and updated independently, and the partial betweenness scores can be summed at the end.

The distributed version of the framework has multiple advantages. First, the space requirements per machine are reduced to $\mathcal{O}(n^2/p)$. For example, if 100 machines are available, for a graph with 1M vertices the storage needed is just 150GB per machine, which is perfectly reasonable by today's standards. Second, the overall work of the algorithm is divided among p processors, leading to a theoretical p speedup over single machine execution. Furthermore, the disk access workload is distributed in a balanced fashion across multiple disks. As a result, the disk access speed of the framework is p times faster. As we show in the next section, the speedup of the framework is indeed almost ideal.

5.3 Online betweenness updates

Real graphs are constantly evolving with addition of new vertices and edges, and removal of existing ones. Updating the betweenness centrality in real-time is extremely challenging, given its computational cost. However, due to the inherent parallelism of our design, the framework can scale not only to large graphs but also to rapidly changing ones.

Each of the p available machines is responsible for updating the data structures and partial centrality scores for n/p sources. The system can monitor the average time $\overline{t_S}$ needed for each machine to process a source, given the addition of a new edge or the removal of an existing one. Thus, the average time $\overline{t_U}$ to produce updated betweenness scores for all vertices and edges upon arrival of an update is $\overline{t_U} = \overline{t_S} \times n/p + \overline{t_M}$, where $\overline{t_M}$ is the average time needed to merge the results.

Assuming that an evolving graph has an average rate of updates per time unit $F = 1/t_I$, the system can always (on average) produce updated betweenness scores before the new time period, if $\overline{t_U} < t_I$. However, if the system measures an increased rate of arrival $F' > F$, it can adjust the number of machines to $p' > \overline{t_S} \times n/(t'_I - \overline{t_M})$ to guarantee online updates, assuming that the average time per source remains unchanged, and that $(t'_I > \overline{t_S} + \overline{t_M})$, i.e., the inter-arrival time is larger than the inherent serial part of the algorithm.

5.4 A MapReduce embodiment

Various paradigms can be used to embody the proposed parallel framework, such as parallel stream processing engines (e.g., Storm, S4, and Samza). Due to its ease of use and popularity, we deploy and experiment with a MapReduce embodiment on a Hadoop cluster.
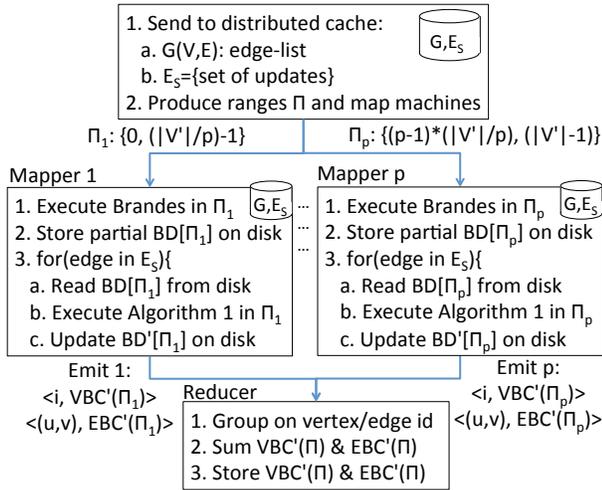
Fig. 4: MapReduce version of our framework.

Figure 4 illustrates a MapReduce adaptation of our algorithmic framework. The graph $G(V, E)$ and the set of updates E_S (new edges to be added and existing edges to be removed) are replicated on all machines via distributed cache, and loaded in memory. We generate an input for each mapper i that represents a partition Π_i of the graph. Each partition is comprised of two integers that represent the first and last ID of the range of sources for which the particular mapper i is responsible. The data structures $BD[\Pi_i]$ created during step 1 are stored locally on the disk of each machine.

The Map function processes all edges in E_S in sequence and updates the betweenness centrality. For each update, it emits key-value pairs of vertex or edge IDs together with their partial betweenness centrality (PBC) by source s, i.e., $\langle id, \mathrm{VBC}_s(id) | \mathrm{EBC}_s(id) \rangle$, where id is either a vertex or an edge identifier. All the intermediate pairs are sent to the reducers who are responsible for producing the final aggregated betweenness results for all vertices and edges. Each Reduce function aggregates the partial betweenness score of one element (vertex or edge) in the graph. The final value of the computation is the new betweenness score for each element of the graph after the set of updates E_S is applied.

The signatures of the functions are as follows:

$$\mathsf{Map}: \quad \langle \Pi_i, \emptyset \rangle \ \rightarrow \ [\langle id, \mathrm{PBC}_s(id) \rangle \ \forall id \in G, \forall s \in \Pi_i]$$
$$\mathsf{Reduce}: \quad \langle id, [\mathrm{PBC}_s(id), \ldots] \ \forall s \in V \rangle \rightarrow \langle id, BC(id) \rangle$$

6 EXPERIMENTAL RESULTS

We evaluate our algorithmic framework on real and synthetic graphs to assess its performance over Brandes' algorithm and its ability to scale to large graphs.

Datasets. We use two types of graphs: synthetic and real. The synthetic graphs are created with a synthetic social graph generator [33], which produces graphs with properties, such as degree distribution and clustering coefficient, similar to real social graphs. The synthetic

TABLE 2: Description of the graphs used. AD: average degree, CC: clustering coefficient, ED: effective diameter.

	Dataset	$|V|(LCC)$	$|E|(LCC)$	AD	CC	ED
synthetic	1k	1000	5895	11.8	0.263	5.47
	10k	10 000	58 539	11.7	0.219	6.56
	100k	100 000	587 970	11.8	0.207	7.07
	1000k	1 000 000	5 896 878	11.8	0.204	7.76
real-world	wikielections	7066	100 780	8.3	0.126	3.78
	slashdot	51 082	117 377	51.1	0.006	5.23
	facebook	63 392	816 885	63.7	0.148	5.62
	epinions	119 130	704 571	12.8	0.081	5.49
	dblp	1 105 171	4 835 099	8.7	0.6483	8.18
	amazon	2 146 057	5 743 145	3.5	0.0004	7.46

graphs enable us to experiment with graphs that maintain properties of real social graphs while being able to freely increase their size (see Table 2).

Table 2 also reports the details of the real graphs we use. They are taken from the KONECT collection[6] and come from different domains: wiki-elections (WE for short, election votes for Wikipedia admins), epinions (EP, trust among Epinion users), facebook (FB, friendships among Facebook users), slashdot (SD, replies among Slashdot users), dblp (co-authorships among scholars), and amazon (AMZ, product ratings by Amazon users). To make the results comparable between real and synthetic graphs as well as with previous works, we use the largest connected component (LCC) of the real graphs.

Graph updates. For edge addition in the synthetic graphs, we generate the stream of added edges E_S by connecting 100 random unconnected pairs of vertices. For the real graphs, each edge has an associated timestamp of its real arrival time, so we simply replay them in order. For edge removal in the synthetic graphs, we randomly select 100 existing edges to construct the stream of removed edges E_S. For real graphs, we remove the last 100 edges that are added in each graph and do not create a graph partition. The use of real arrival times in graphs is an important difference from previous studies of betweenness centrality updates. This scenario allows us to simulate the evolution of a real system, and thus assess the capability of our framework to update the betweenness centrality *online*.

Implementation. We implement our algorithmic framework in Java and use the JUNG graph library[7] for basic graph operations and maintenance. For the out-of-core version, we store $BD[\cdot]$ in a single file, and read it sequentially in memory source by source. If any update is needed for the current source, it is performed in place on disk rather than overwriting the whole file. This enhancement limits the writes on disk to a minimum. In the experiments, we compare the performance of three versions of the framework: (1) in memory with prede-
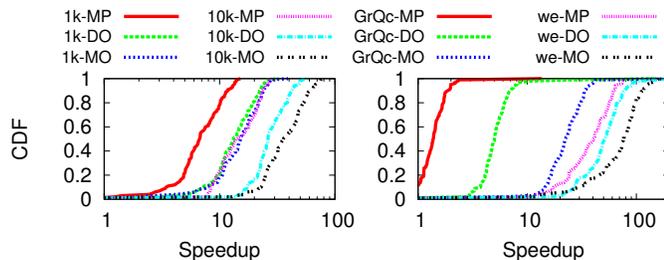
---

6. http://konect.uni-koblenz.de/graphs
7. http://jung.sourceforge.net

Fig. 5: Speedup of the framework's 3 versions on synthetic and real graphs executed on single machines (addition).

cessors lists (MP), (2) in memory without predecessors lists (MO), (3) on disk without predecessors lists (DO).

Infrastructure. For the single-machine version (both in-memory and out-of-core), we use high-end servers with 8-core Intel Xeon @2.4GHz CPU and 50GB of RAM each. For the parallel version, we use a Hadoop cluster with hundreds of machines with 8-core Intel Xeon @2.4GHz CPU and 24GB of RAM each. We report the average performance over 10 executions of the algorithms for each experimental setup.

6.1 Speedup over Brandes' algorithm

Predecessors list optimization. Figure 5 presents the cumulative distribution function (CDF) of the speedup over Brandes' algorithm when testing the three versions of the framework on edge addition. Each point in the graph represents the speedup when adding one of the 100 edges in the stream (averaged over 10 runs). The results show that removing the predecessors lists can actually boost the performance (the MO version is always faster than the MP version). As the algorithm does not need to create the lists nor to maintain them, the overhead, and thus the overall execution time, is reduced.

Related work comparison. As shown in Table 3, the average performance of our framework is comparable to the reported results of previously proposed techniques. The method in [21] shows faster results in networks with low clustering coefficient where changes do not affect many vertices. However, in social networks with high clustering, the method by Kas et al. [21] and especially QUBE [25] are slow due to high cost in updating vertex centrality (in QUBE, many vertices might be included in the minimum union cycle to be updated).

We also compared with the method by Green et al. [17] using a version of their code. By taking into account implementation differences (Java vs. C), the speedups observed are comparable to our method. However, under limited main memory, the method by Green et al. fails to tackle a medium-sized network like slashdot. Instead, as demonstrated later in Table 4, our framework can handle even larger datasets using out-of-core techniques with small main memory footprint and significant speedups.

TABLE 3: Speedup comparison with related work.

Dataset	$|V|$	MO avg (max)	[21]	[25]	[17]
wikivote	7k	75 (181)		3	
contact	10k	75 (153)		4	
UCI (fb-like)	2k	32 (90)	18		
ca-GrQc	4k	31 (378)	68	2	40
ca-HepTh	8k	42 (80)	358		40
adjnoun	.1k	48 (172)			20
ca-CondMat	19k	94 (395)			109
as-22july06	23k	70 (291)			61
slashdot (50GB)	51k	88 (178)			X

Additionally, these past techniques compute only vertex centrality, while our method computes both vertex and edge centrality with the shown speedups.

Out-of-core performance. When $BD[\cdot]$ is stored on disk (DO) rather than in memory (MO), we observe a decrease of the speedup due to the slower access time of the disk. Overall, the DO version is more than $10\times$ faster than Brandes' for the 1k and more than $30\times$ for the 10k graph (median values). The time to process a single edge depends heavily on which parts of the graph it connects and how many structural changes it produces. The in-memory version is CPU bound, so this variability is reflected in the execution time. On the other hand, the out-of-core version is I/O bound, and the execution time is dominated by disk access and the variability of the CPU time spent becomes latent. In the remainder of this section we use the DO version. Key speedup results are summarized in Table 4.

MapReduce speedup. Figure 6(a) shows the CDF of speedup over Brandes' algorithm when executing the DO version on a MapReduce cluster for addition of edges. In the experiment, we adjust the number of mappers so that each mapper is assigned 1k sources per graph. Brandes' algorithm is compared with the cumulative execution time of our algorithm, i.e., the sum of execution times across all mappers and reducers. By increasing the graph size from 1k to 100k vertices, the median speedup increases from ≈ 10 to ≈ 50. When increasing the graph size to 1000k vertices, the median speedup drops to ≈ 10. Compared to the experiments on a single machine, there

TABLE 4: Summary of key speedup results.

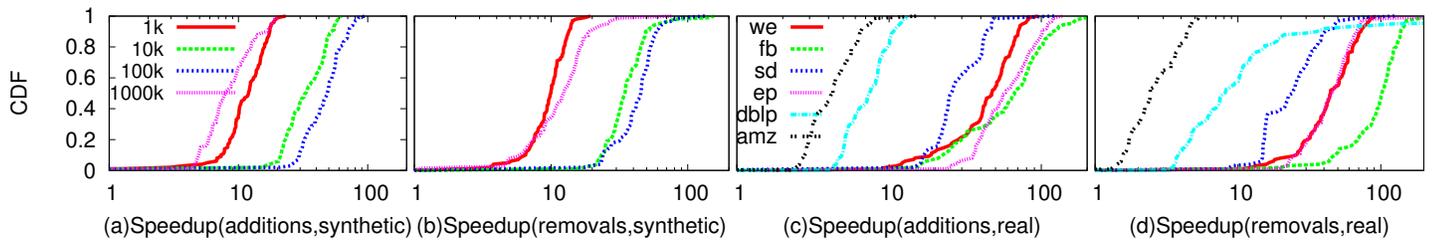| Dataset	Addition			Removal		
	Min	Med	Max	Min	Med	Max
1k	3	12	23	2	10	19
10k	16	34	62	2	35	155
100k	21	49	96	4	45	134
1000k	5	10	20	1	12	78
wikielections	9	47	95	1	45	92
slashdot	15	25	121	8	24	127
facebook	10	66	462	1	102	243
epinions	24	56	138	2	45	90
dblp	3	8	15	3	8	429
amazon	2	4	15	2	3	5

Fig. 6: Speedup of DO version on synthetic/real graphs executed on a MapReduce cluster (additions/removals).

is an increase in the variability of the framework's performance when accessing the disks on the MapReduce cluster. This effect is partly due to contention on disk access on the cluster caused by concurrent jobs, as well as increased computation load per machine and source. Overall, the use of parallel execution leads to improved speedups for larger graphs that would be impossible to process on a single machine. As an additional benefit, we also get reduced wall-clock time.

Figure 6(b) shows the CDF of speedup over Brandes' algorithm on a MapReduce cluster for removal of edges. The setup is similar to the previous experiment. By increasing the graph size from 1k to 100k vertices, the median speedup increases from ≈ 10 to ≈ 45. When increasing the graph size to 1000k vertices, the median speedup drops to ≈ 12. In this case, the speedup is slightly higher than when adding edges, because the removal of edges reduces the shortest paths between vertices and causes slightly less computational load.

Real graph structure. Figures 6(c) and (d) show the CDF of speedup over Brandes' algorithm for the real graphs when adding or removing edges, respectively. Also in this case we adjust the number of mappers so that each mapper is assigned 1k sources per graph.

In the edge addition, facebook exhibits the highest variability with a median speedup of ≈ 66. In the edge removal, dblp exhibits higher variability than facebook with a median speedup of ≈ 8. When adding edges on slashdot, which has a number of edges similar to wikielections but fewer vertices, our framework exhibits lower variability and smaller maximum speedup than on wikielections. It also performs better on facebook than slashdot, both in addition and removal of edges, even if the two graphs have approximately the same number of vertices. One reason may lie in the higher clustering coefficient of wikielections and facebook, which reduces the number of structural changes upon update.

In support to our hypothesis, for amazon we observe a low median speedup of ≈ 4. The low performance is due to the structural properties of this graph: very low clustering coefficient and high diameter lead to many structural changes upon edge addition or removal, and thus higher computational load. For example, in both addition and removal of edges, we observe that on dblp, which is of the same order of magnitude as amazon but with a much higher clustering coefficient, our method achieves about double the speedup than
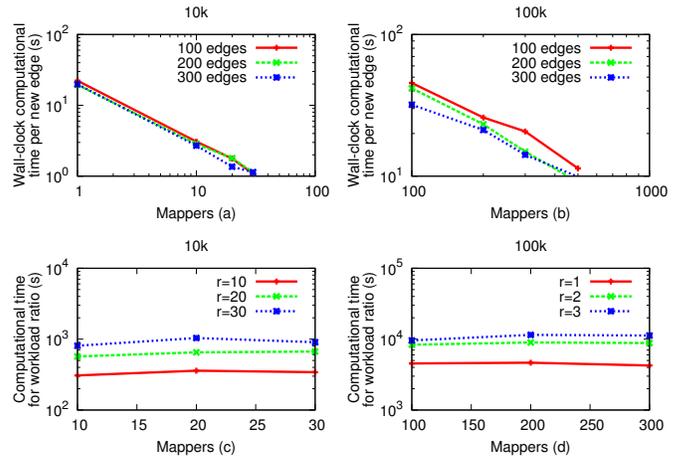

Fig. 7: (a-b) Computation time under increasing number of mappers. (c-d) Computation time under constant ratio of workload over mappers.
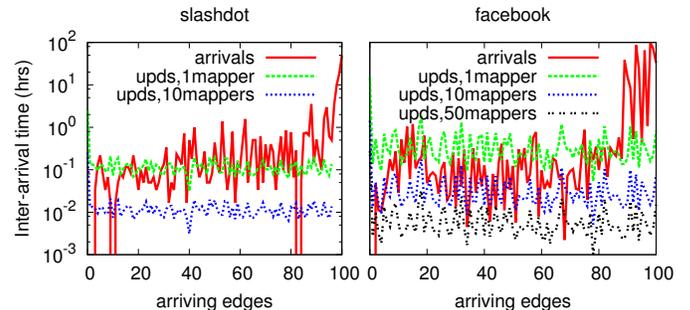

Fig. 8: Inter-arrival time of facebook and slashdot edges and update times for betweenness centrality.

on amazon. We conjecture that performance on a larger graph is tightly connected with its structural properties, the longer disk access time, and more computational load per source (more vertices to be traversed). Exploring the connections between algorithm's performance and graph properties is an interesting path for future investigation.

6.2 Scalability for online updates

Figures 7(a-b) analyze the strong scaling properties of the algorithm in the case of edge addition. In these experiments we keep the workload fixed and increase the parallelism level. By employing a larger number of

TABLE 5: Edges missed and average delay vs. scaling.

Dataset	mappers	% missed	avg. delay (s)
slashdot	1	44.565	257.9
slashdot	10	1.087	32.4
facebook	1	69.697	1061.1
facebook	10	19.192	96.6
facebook	50	3.030	8.6
facebook	100	1.010	5.5

Fig. 9: Girvan-Newman: continuous removal of edge with highest betweenness and re-computation on synthetic graphs.

mappers, the overall execution time decreases almost linearly regardless of the workload (i.e., 100 vs. 300 added edges) and graph size (i.e., 10k vs. 100k). As expected, our algorithm is embarrassingly parallel and shows very good scalability properties.

Figures 7(c-d) explore the weak scaling properties of the algorithm in the case of edge addition. In these experiments we keep the workload per processing unit fixed, and we increase the workload and parallelism level proportionally. The workload for the system is represented by the number of edges updates within a period of time. To keep the ratio of workload per mapper constant, we increase the number of edges updates proportionally to the number of mappers. As shown in the figures, the total execution time remains constant at different levels of ratio (e.g., 1 vs 3, or 10 vs. 30) regardless of the parallelism level. These results show that our parallel algorithmic framework can scale to larger workload simply by adding more machines.

Figure 8 demonstrates the *online* capabilities of the algorithm on two real graphs. The figure shows the inter-arrival time of new edges and the time needed by the framework to produce updated betweenness values. In Table 5 we report the fraction of edges for which the framework was unable to produce updates on time and the corresponding average delay. For slashdot, the framework manages to produce online updates for 98.91% of edges with 10 mappers. For facebook instead, the arrival rate is higher and 10 machines are not enough. However, the scalability of our algorithm allows to simply use more machines to decrease the response time of the system. Thus, with 100 mappers, the system can produce online updates for 98.99% of edges.

6.3 Use-case: Girvan-Newman community detection

Our framework allows a faster update of the betweenness of edges, thus enabling several applications including community detection with the Girvan-Newman method [15]. This method relies primarily on the computation of the betweenness of all edges. It iteratively removes the edge with the highest centrality to create disconnected components. The procedure is repeated until no edges remain, thus enabling the construction of a hierarchy of communities. In practice, this algorithm has been abandoned due to the high cost incurred from the recalculation of the betweenness on the modified graph. However, our framework allows a faster update of betweenness by taking into account only the affected
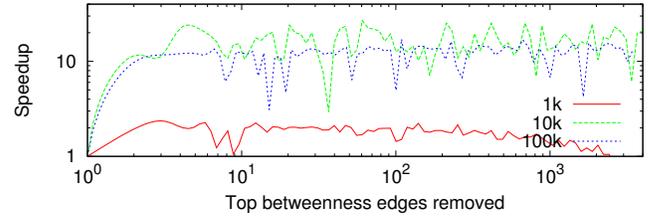
parts of the graph, thus attaining an order of magnitude faster execution over Brandes when running the Girvan-Newman algorithm on a single machine (Figure 9).

7 CONCLUSIONS

The computational complexity of most existing graph algorithms makes them impractical in nowadays massive and dynamic networks. In order to scale graph analysis to real-world applications and to keep up with their highly dynamic nature, we need to devise new approaches specifically tailored for modern parallel stream processing engines that run on clusters of shared-nothing commodity hardware.

In this paper we introduce an algorithmic framework for computing betweenness centrality incrementally in large evolving graphs where edges and vertices are added and removed. Our experimental results demonstrate that our framework is capable of scaling-out to a large number of machines without additional overhead. This feature allows it to process graphs whose size is orders of magnitude larger than what previously reported in the literature.

Moreover, our method shows very good scale-up properties, which lead to an almost linear decrease of the execution time needed to update the betweenness centrality on parallel systems. As a result, our method is able to keep up with the incoming rate of updates in large real-world graphs and in an online fashion.

The scalability achieved by our framework opens the doors to new applications for real-world networks. For instance, our framework can be exploited for online detection and prediction of emerging leaders and communities in social networks.

8 REFERENCES

[1] Sariyüce, Ahmet Erdem and Kaya, Kamer and Saule, Erik and Çatalyürek, Ümit V. Betweenness Centrality on GPUs and Heterogeneous Architectures. In *6th Workshop on General Purpose Processor Using Graphics Processing Units*, pages 76–85, 2013.

[2] C. S. Ang. Interaction networks and patterns of guild community in massively multiplayer online games. *Social Network Analysis and Mining*, 1(4):341–353, 2011.

[3] J. Anthonisse. The rush in a directed graph. Technical Report BN9/71, Stichting Mathematisch Centrum, Amsterdam, Netherlands, 1971.

[4] D. Bader and K. Madduri. Parallel algorithms for evaluating centrality indices in real-world networks. In *Int. Conf. on Parallel Processing*, pages 539–550, 2006.

[5] P. Boldi and S. Vigna. Axioms for centrality. http://arxiv.org/abs/1308.2140, 2013.

[6] U. Brandes. A faster algorithm for betweenness centrality. *J. of Mathematical Sociology*, 25(2):163–177, 2001.

[7] U. Brandes. On variants of shortest-path betweenness centrality and their generic computation. *Social Networks*, 30(2):136–145, May 2008.

[8] U. Brandes and C. Pich. Centrality estimation in large networks. *Int. J. of Bifurcation and Chaos*, 17(7):2303–2318, 2007.

[9] R. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, 1992.

[10] S. Catanese, E. Ferrara, and G. Fiumara. Forensic analysis of phone call networks. *Social Network Analysis and Mining*, 3(1):15–33, 2012.

[11] G. De Francisci Morales. SAMOA: A Platform for Mining Big Data Streams. In *Int. Work. on Real-Time Analysis & Mining of Social Streams@WWW*, pages 777–778, 2013.

[12] G. De Francisci Morales and A. Bifet. SAMOA: Scalable Advanced Massive Online Analysis. *JMLR*, 16(Jan): 149–153, 2015.

[13] C. Freeman, S. Borgatti, and D. White. Centrality in valued graphs: A measure of betweenness based on network flow. *Social Networks*, 13(2):141–154, 1991.

[14] L. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.

[15] M. Girvan and M. E. J. Newman. Community structure in social and biological networks. *National Academy of Sciences of USA*, 99(12):7821–7826, June 2002.

[16] M. Granovetter. The strength of weak ties. *Amer. J. of Sociology*, 78(6):1360–80, May 1973.

[17] O. Green, R. McColl, and D. A. Bader. A fast algorithm for streaming betweenness centrality. In *SOCIALCOM-PASSAT*, pages 11–20, 2012.

[18] O. Green, and D. A. Bader. Faster betweenness centrality based on data structure experimentation. In *International Conference on Computational Science*, 2013.

[19] H. Jeong, S. Mason, A. Barabási, and Z. Oltvai. Lethality and centrality in protein networks. *Nature*, 411:41, 2001.

[20] G. Kahng, E. Oh, B. Kahng, and D. Kim. Betweenness centrality correlation in social networks. *Phys. Rev. E*, 67: 017101, 2003.

[21] M. Kas, M. Wachs, K. M. Carley, and L. R. Carley. Incremental algorithm for updating betweenness centrality in dynamically growing networks. In *ASONAM*, pages 33–40, August 2013.

[22] N. Kourtellis and A. Iamnitchi. Leveraging peer centrality in the design of socially-informed peer-to-peer systems. *IEEE Transactions in Parallel and Distributed Systems*, 99:1, 2013.

[23] N. Kourtellis, T. Alahakoon, R. Simha, A. Iamnitchi, and R. Tripathi. Identifying high betweenness centrality nodes in large social networks. *Social Network Analysis and Mining*, 3:899–914, 2013.

[24] N. Kourtellis, G. De Francisci Morales, and F. Bonchi. Scalable online betweenness centrality in evolving graphs. http://arxiv.org/abs/1401.6981, 2014.

[25] M.-J. Lee, J. Lee, J. Park, R. Choi, and C.-W. Chung. A quick algorithm for updating betweenness centrality. In *Int. Conf. on World Wide Web*, pages 351–360, April 2012.

[26] F. Liljeros, C. Edling, L. Amaral, H. Stanley, and Y. Aberg. The web of human sexual contacts. *Nature*, 411:907, 2001.

[27] A. McLaughlin and D. Bader. Revisiting Edge and Node Parallelism for Dynamic GPU Graph Analytics. In Workshop on Multithreaded Architectures and Applications (MTAAP), 2014.

[28] L. A. Maglaras and D. Katsaros. New measures for characterizing the significance of nodes in wireless ad hoc networks via localized path-based neighborhood analysis. *Social Network Analysis and Mining*, 2:97–106, 2012.

[29] M. Nasre, M. Pontecorvi, and V. Ramachandran. Betweenness Centrality Incremental and Faster. In *Mathematical Foundations of Computer Science*, pages 577–588, LNCS(8635), 2014.

[30] M. Newman. A measure of betweenness centrality based on random walks. *Social Networks*, 27(1):39–54, 2005.

[31] G. Ramalingam and T. Reps. An Incremental Algorithm for a Generalization of the Shortest-Path Problem. *Journal of Algorithms*, 21(2):267–305, 1996.

[32] M. Riondato and E. M. Kornaropoulos. Fast Approximation of Betweenness Centrality Through Sampling. In *WSDM*, 2014.

[33] A. Sala, L. Cao, C. Wilson, R. Zablit, H. Zheng, and B. Y. Zhao. Measurement-calibrated graph models for social network experiments. In *19th Int. Conf. on the World Wide Web*, pages 861–870, Raleigh, NC, USA, 2010.

Nicolas Kourtellis is a Postdoctoral Researcher in the Web Mining Research Group at Yahoo Labs Barcelona. He received a PhD in Computer Science and Engineering from the University of South Florida in 2012. He is interested in the network analysis of large-scale systems and social graphs and property extraction useful in the design of improved socially-aware distributed systems. More information at http://www.csee.usf.edu/~nkourtel/.

Gianmarco De Francisci Morales is a Visiting Scientist at Aalto University, Helsinki, Finland. He is an active member of the Apache Software Foundation working on the Hadoop ecosystem (Giraph, S4), and a committer for the Apache Pig project. He co-organizes the workshop series on Social News on the Web (SNOW) co-located with the WWW conference. He is a lead developer of Apache SAMOA, a platform for mining big data streams. More information at http://gdfm.me.

Francesco Bonchi is a Director of Research at Yahoo Labs Barcelona, where he is the head of the Web Mining Research Group. He is member of the ECML PKDD Steering Committee and he serves in the editorial board of ACM Transactions on Intelligent Systems and Technology (TIST), and IEEE Transactions on Knowledge and Data Engineering (TKDE). More information at http://www.francescobonchi.com/.