



Approximate 2-hop neighborhoods on incremental graphs: An efficient lazy approach

Luca Becchetti
Sapienza University of Rome
Rome, Italy
becchetti@diag.uniroma1.it

Andrea Clementi
Tor Vergata University of Rome
Rome, Italy
clementi@mat.uniroma2.it

Luciano Gualà
Tor Vergata University of Rome
Rome, Italy
guala@mat.uniroma2.it

Luca Pepè Sciarria
Tor Vergata University of Rome
Rome, Italy
luca.pepesciarria@gmail.com

Alessandro Straziota
Tor Vergata University of Rome
Rome, Italy
alessandro.straziota@uniroma2.it

Matteo Stromieri
Tor Vergata University of Rome
Rome, Italy
matteo.stromieri@students.uniroma2.eu

ABSTRACT

In this work, we propose, analyze and empirically validate a lazy-update approach to maintain accurate approximations of the 2-hop neighborhoods of dynamic graphs resulting from sequences of edge insertions.

We first show that under random input sequences, our algorithm exhibits an optimal trade-off between accuracy and insertion cost: it only performs $O(\frac{1}{\varepsilon})$ (amortized) updates per edge insertion, while the estimated size of any vertex's 2-hop neighborhood is at most a factor ε away from its true value in most cases, *regardless* of the underlying graph topology and for any $\varepsilon > 0$.

As a further theoretical contribution, we explore adversarial scenarios that can force our approach into a worst-case behavior at any given time t of interest. We show that while worst-case input sequences do exist, a necessary condition for them to occur is that the *girth* of the graph released up to time t be at most 4.

Finally, we conduct extensive experiments on a collection of real, incremental social networks of different sizes, which typically have low girth. Empirical results are consistent with and typically better than our theoretical analysis anticipates. This further supports the robustness of our theoretical findings: forcing our algorithm into a worst-case behavior not only requires topologies characterized by a low girth, but also carefully crafted input sequences that are unlikely to occur in practice.

Combined with standard sketching techniques, our lazy approach proves an effective and efficient tool to support key neighborhood queries on large, incremental graphs, including neighborhood size, Jaccard similarity between neighborhoods and, in general, functions of the union and/or intersection of 2-hop neighborhoods.

PVLDB Reference Format:

Luca Becchetti, Andrea Clementi, Luciano Gualà, Luca Pepè Sciarria, Alessandro Straziota, and Matteo Stromieri. Approximate 2-hop neighborhoods on incremental graphs:
An efficient lazy approach. PVLDB, 18(11): 3937 - 3950, 2025.
doi:10.14778/3749646.3749665

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 18, No. 11 ISSN 2150-8097.
doi:10.14778/3749646.3749665

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/Gnumlab/graph_ball.

1 INTRODUCTION

In this paper, we consider the task of processing a possibly large, dynamic graph $G(V, E)$, incrementally provided as a stream of edge insertions, so that at any point of the stream it is possible to efficiently evaluate different queries that involve functions of the *h -hop neighborhoods* of its vertices. For a vertex $v \in V$, its *h -hop neighborhood* is simply the *set* of vertices that are within h hops from v . In the remainder, *h -hop neighborhoods* are called *h -balls* for brevity. As concrete examples of query types we consider, one might want to estimate the size of the 2-ball at any given vertex, or the Jaccard similarity between the 2-balls centered at any given two vertices, or other indices of a similar flavor that depend on the intersection or union between 1-balls and/or 2-balls, just to mention a few.

Neighborhood-based indices are common in key mining tasks, such as link prediction in social [40] and biological networks [52] or to describe statistical properties of large social graphs [6]. For example, 2-hop neighborhoods are important in social network analysis and similarity-based link prediction [4, 55, 57], while accurate approximations of *h -balls'* sizes are used to estimate key statistical properties of (very) large social networks [3, 10], or as link-based features in classifiers for Web spam detection [7]. Further tasks that may explicitly involve or benefit from 2-hop neighborhood queries are accurate estimates of centrality measures that are widely adopted in social network analysis [9, 29, 46] or mining of (usually large) bipartite networks that are frequent in certain applications [43, 49]. We investigate the former example in detail in the full version [8].

When the graph is static, an effective approach to this general task is to treat *h -balls* as subsets of the vertices of the graph, suitably represented using approximate summaries or sketches [1]. This line of attack has proved successful, for example in the efficient and scalable evaluation of important neighborhood-based queries on massive graphs that in part or mostly reside on secondary storage [6, 10, 27, 41].

Nowadays, standard applications in social network analysis often entail dynamic scenarios in which input graphs *evolve over time*,

under a sequence of edge insertions and possibly deletions [2]. In many cases, such as co-authorship, citation or user-item association networks, the corresponding dynamic graphs are inherently or mostly *incremental*, i.e., either edges and/or vertices are added but not deleted, or insertions account for the vast majority of updates [12, 38].¹

With respect to a static scenario, the dynamic case poses new and significant challenges even in the incremental setting, as soon as $h > 1$.² To see this, it may be useful to briefly sketch the cost of maintaining 1- and 2-balls exactly under a sequence of edge insertions, as we discuss in more detail in Section 2. When $h = 2$, each $\text{Insert}(u, v)$ operation entails (see Algorithm 1 and Figure 1): i) updating the 2-ball of u to its union with the 1-ball of v and viceversa (what we call a *heavy* update); ii) adding v to the 2-ball of *each* neighbor of u and viceversa (what we call a *light* update). Both heavy and light updates can result in high computational costs per edge insertion: a heavy update can be expensive if at least one of the neighborhoods to merge is large; on the other hand, light updates are relatively inexpensive, but they can be numerous when large neighborhoods are involved, again resulting in a high overall cost per edge insertion. Unfortunately, h -balls can grow extremely fast with h in many social networks, already as one switches from $h = 1$ to $h = 2$ [3, 7]. For the same reason, maintaining lossless representations of 2-balls for each vertex of such networks might require considerable memory resources and might negatively impact the cost of serving neighborhood-based queries that involve moderately or highly central vertices.

To address the aforementioned issues for graphs that reside in main memory, one might want to trade some degree of accuracy for the following broad goals: 1) designing algorithms with low update costs, possibly $O(1)$ amortized per edge insertion; 2) minimizing memory footprint beyond what is needed to store the graph; 3) maintaining 1- and 2-balls using data structures that afford efficient, real-time computation of queries as the ones mentioned earlier with minimal memory footprint.

Heavy updates are natural and well-known candidates for efficient (albeit approximate) implementation using compact, sketch-based data structures [1, 5, 14, 15, 30]. However, sketches alone are of no avail in handling light updates, whose sheer potential number requires a novel approach. The literature on efficient data structures that handle insertions and often deletions over dynamic graphs is rich. However, efficient solutions to implement neighborhood-based queries on dynamic edge streams are only known for 1-balls [17, 20, 33, 53], nor do approaches devised for other dynamic problems adapt to our setting in any obvious way, something we elaborate more upon in Section 1.2.

1.1 Our Contribution

In this paper, we propose an approach that trades some degree of accuracy for a substantial improvement in the average number of

light updates. In a nutshell, upon an edge insertion, our algorithm performs the (two) corresponding heavy updates, but in general only a subset of the required light updates, according to a scheme that combines a threshold-based mechanism and a randomized, batch-update policy. Hence, for every vertex u , we only keep an approximation (a subset to be specific) of u 's 2-ball. If 1- and 2-balls are represented with suitable data sketches, our approach affords constant average update cost per edge insertion.³ While the behavior and accuracy guarantees of most sketching techniques are well understood, the estimation error induced by lazy updates can be arbitrarily high in some cases. The main focus of this paper is on the latter aspect, which is absent in the static case but critical in the dynamic setting. Accordingly, we assume lossless representations of 1-balls and approximate 2-balls in our theoretical analyses in Sections 2 to 4, while we use standard sketching techniques to represent 1- and 2-balls in the actual implementations of the algorithms and baselines we consider in the experimental analysis discussed in Section 5.

Almost-optimal performance on random sequences. We prove in Section 3 that even a simplified, deterministic variant of our lazy-update Algorithm 3 achieves asymptotically optimal expected performance when the sequence of edge insertions is a random, uniform permutation over an *arbitrary* set of edges. In other words, our lazy approach is robust to adversarial topologies as long as the edge sequence follows a random order. Formally, we prove that, for any desired $0 < \epsilon < 1$, our algorithm only performs $O(\frac{1}{\epsilon})$ (amortized) updates per edge insertion, while at any time t and for every vertex v , the estimated size of v 's 2-ball is, in expectation, at most a factor ϵ away from its true value. We further prove that this approximation result holds with a probability that exponentially increases with the true size of the 2-ball itself (Theorem 1). Thanks to this analysis in concentration, our results can be extended to other functions of 2-balls, including union, intersection and Jaccard similarity (see Corollary 2 for this less obvious case).

As positive as this result may sound, it begs the following questions from a careful reader: 1) Are the results above robust to adversarial sequences? 2) Is a performance analysis under random sequences representative of practical scenarios? 3) More generally, does our lazy scheme offer significant practical advantages?

Performance analysis on adversarial inputs. While our results for random sequences are optimal regardless of the underlying graph's topology, one might wonder about the ability of an adversary to design *worst-case, adaptive sequences* that force our approach to behave poorly and, in this case, whether any conditions on the graph topology are *necessary* for this to happen. We investigate this issues in Section 4, where we first show that it is possible to design worst-case sequences of edge insertions that force our algorithm to perform arbitrarily worse than the random setting (Theorem 3). However, as a further contribution, we also prove that worst-case

¹Of course, notable examples exist in which deletions are as important, such as some affiliation networks, or networks that describe on-line messaging behavior, where edges representing message exchanges become less and less important as they age.

²The case $h = 1$ is considerably simpler and it *barely relates to graphs*: adding or removing one edge (u, v) simply requires updating the 1-balls of u and v accordingly, i.e., updating two corresponding set sketches by adding or removing one item. This has been the focus of extensive work in the recent past that we discuss in Section 1.2.

³The particular sketch used depends on the neighborhood queries we want to be able to serve. When sketches are used, the cost of merging two neighborhoods corresponds to the cost of combining the corresponding sketches, which is typically a constant that depends on the desired approximation guarantees. For example, if we are interested in the Jaccard similarity between pairs of 1- and/or 2-balls, this cost will be proportional to the (constant) number of minhash values we use to represent each neighborhood.

input sequences exist *only if* the *girth* [22] of the final graph is at most 4. More precisely, we show that a randomized, special case of Algorithm 3 achieves asymptotically optimal performance on a class of graphs that contains all graphs with girth at least 5,⁴ even when the input sequence is chosen by an adaptive adversary.

Experimental analysis. An analysis under random permutation sequences is relatively common in the literature on dynamic edge streams and data streams [16, 32, 34, 44]. Yet, one might reasonably wonder about its practical significance for the task considered in this paper. We address these issues in Section 5, where we conduct experiments on small, medium and large-sized, incremental graphs (whose main properties are summarized in Table 1). At least on the diverse sample of real networks we consider, experimental results on the estimation of key queries such as size and Jaccard similarity are consistent with the theoretical findings from Section 3, while results on running times highlight the significance of our approach in practice: i) light updates crucially affect computational costs in practical scenarios; ii) when combined with sketches, our approach strikes an excellent balance between loss in accuracy and computational efficiency. In particular, its accuracy is comparable or only slightly worse than the baseline that performs all necessary light updates, while it achieves speed-ups that increase with network size and can be orders of magnitude with respect to the baseline on large or very large social networks (Tables 4 and 5). We finally note that our datasets are samples of real social networks. As such, they have relatively large local and global clustering coefficients⁵ and thus low girth. Hence, our experimental analysis further supports the robustness of our theoretical findings: forcing our algorithm(s) into a worst-case behavior not only requires topologies characterized by a low girth, but also carefully crafted input sequences that are unlikely to occur in practice.

Use cases from real networks. While neighborhood-based queries are extremely important in social network analysis, the main focus of this paper is on crucial, new challenges that arise in dynamic settings and on possible strategies to address them. For this reason, we investigate basic neighborhood queries without specific applications in mind. As simple as it may appear however, the incremental setting we consider provides a feasible framework to tackle important mining tasks over large, dynamic networks. As an example, we consider the problem of maintaining the *harmonic centrality*-based ranking of the vertices of an incremental graph [11, 46], a key mining task in social network analysis [11, 29]. Exact solutions for this problem necessarily incur high computational costs in some cases,⁶ while the study of dynamic strategies with update costs of practical interest is still an open area of research [31]. In the full version of this paper [8], we experimentally show that our approach is a very good fit for this task, affording continuous, accurate tracking of top- k ranking vertices, thus highlighting its effectiveness for a key downstream task that would otherwise be hardly feasible on large, incremental graphs.

Many real networks are inherently bipartite, describing interactions between entities of two different types. Examples include

many biological networks [43], social networks [35], user-URL or user-item interactions [37] to name a few. Moreover, many of them are also inherently or mostly incremental in nature.⁷ In such networks, the prediction of future interactions is a key mining task, for which (Jaccard) similarity between 2-hop neighborhoods can prove a valuable index for prediction purposes. For example, in author-paper networks, the set $L_2(v)$ of vertices at distance exactly 2 from an author v is the set of v 's co-authors. Hence, for any two authors u and v with no currently co-authored paper, the Jaccard similarity between $L_2(u)$ and $L_2(v)$ can be used to estimate the probability of u and v co-authoring a paper in the next future.⁸

Final remarks. We stress that although we present them for the undirected case for ease of presentation and for the sake of brevity, our algorithms apply to directed graphs in general,⁹ while our analysis extends to the directed case with minor modifications. Moreover, our algorithms seamlessly address incremental settings in which new vertices may connect to the current graph, the latter being for example the typical scenario in co-authorship networks [12].

1.2 Further related work

Efficient data structures for queries that involve h -balls of a dynamic graph turn out to be useful in different network applications, as we mentioned earlier [4, 7, 55].

Efficient solutions for Jaccard similarity queries on 1-balls have been proposed for different dynamic graph models: all of them share the use of suitable data sketches to manage insertion and deletion of elements from sets. In particular, [17, 33, 53] proposes and compares different approaches that work in the fully-dynamic streaming model, while an efficient solution, based on a buffered version of the k -min-hashing scheme is proposed in [20]. A further algorithm is presented in [56], where *bottom- k sketches* [21] are used to perform dynamic graph clustering based on Jaccard similarity among vertices' neighborhoods. We remark that none of these previous approaches include ideas or tools that can be adapted to efficiently manage the 2-ball update-operations we need to implement in this work.

As for other queries that might be "related" or "useful" in our setting, a considerable amount of work on data structures that support edge insertions and deletions exists for several queries, such as connectivity or reachability, (exact or approximate) distances, minimum spanning tree, (approximate) *betweenness centrality*, and so on. We refer the reader to [31] for a nice survey on experimental and theoretical results on the topic. To the best of our knowledge however, none of these approaches can be obviously adapted to handle the types of queries we consider in this work. For example, a natural idea would be using an incremental data structure to dynamically maintain the first h levels of a BFS tree, such as [26, 47], that achieve $O(h)$ amortized update time. However, let alone effectiveness in efficiently serving queries as the ones we consider here, the data structure uses $\Omega(n)$ space per BFS. This is prohibitive

⁷Edge/vertex deletions are rare or absent.

⁸Our approach also allows estimation of the Jaccard similarity between $L_2(u)$ and $L_2(v)$.

⁹Of course, in this case we have directed h -balls, i.e., sets of vertices that can be reached in at most h hops from a given vertex, or from which it is possible to reach the vertex under consideration in at most h hops.

⁴The class is in fact more general since it also includes graphs with a "bounded" number of cycles of length at most 4. See Definition 4.1, for a formal definition of this class.

⁵At least the undirected ones.

⁶Even in the static setting [9].

in our setting, where we would need to instantiate one such data structure for each vertex, with total space $\Omega(n^2)$. Moreover, since in a degree- Δ graph $\Theta(\Delta)$ BFS trees can change following a single edge insertion, the corresponding amortized time per edge insertion could be as high as $\Theta(\Delta)$, which is basically the same cost of the baseline solution we discuss at the beginning of Section 2.

2 LAZY-UPDATE ALGORITHMS

Preliminaries and notations. The dynamic (incremental) graph model we study can be defined as a sequence $\mathcal{G} = \{G^{(0)}(V, E^{(0)}), \dots, G^{(t)}(V, E^{(t)}), \dots, G^{(T)}(V, E^{(T)})\}$, where: (i) the set of vertices $V = \{1, \dots, n\}$ is fixed, (ii) $T \leq \binom{n}{2}$ is the final graph, while (iii) $E^{(t)}$ is the subset of edges at time t . Note that this changes in every time step $t \geq 1$, as a new edge $e^{(t)}$ is inserted, so that $E^{(t+1)} = E^{(t)} \cup \{e^{(t)}\}$. We remark that our analysis and all our results can be easily adapted to a more general model that includes any combination of the following variants: (i) growing vertex sets, (ii) multiple insertions of the same edge, and (iii) directed edges (thus yielding directed graphs). However, the corresponding adaptations of our analysis would require significantly heavier notation and some technicalities that we decided to avoid for the sake of clarity and space.

Our goal is to design algorithms that, at every time step $t \geq 1$, are able to efficiently compute queries over the current 2-balls of $G^{(t)}$. As mentioned in the introduction, our focus is on queries that are typical in graph mining such as: (i) given a vertex u , estimate the size of $B_2(u)$, and (ii) given two vertices $u, v \in V$, estimate the Jaccard similarity of the corresponding 2-balls:

$$J(B_2(u), B_2(v)) = \frac{|B_2(u) \cap B_2(v)|}{|B_2(u) \cup B_2(v)|}.$$

Both the theoretical and experimental analysis of our lazy-update algorithms consider the following key performance measures: the *amortized update time* per edge insertion and the *approximation ratio* of our algorithms on the quantities $|B_2(u)|$ and $J(B_2(u), B_2(v))$, for any choice of the input vertices. Intuitively, the amortized update time is the average time it takes to process a new edge, a more formal definition is deferred to Section 2, after a detailed description of the algorithms we consider.

We next summarize notation that is used in the remainder of the paper. For a vertex $v \in V$ of a graph $G(V, E)$, we define:

- $\mathcal{N}(v)$: the set of neighborhoods of the vertex v .
- \deg_v : the degree of v . Notice that $\deg_v = |\mathcal{N}(v)|$;
- $L_h(v)$: set of vertices at distance exactly h from v ;
- $B_h(v)$: set of vertices at distance at most h from v .

The reader may have noticed that, in our notation above, the term t does not appear: this is due to the fact that our analysis holds at any (arbitrarily fixed) time step, which is always clear from context.

2.1 Algorithm description

Consider the addition of a new edge (u, v) to G . Clearly, the only 2-balls that are affected are those centered at u , v , and at every vertex $w \in \mathcal{N}(u) \cup \mathcal{N}(v)$. A *baseline* strategy, given as Algorithm 1 for the sake of reference, tracks changes exactly and thus updates all 2-balls that are affected by an edge insertion.

Algorithm 1: Baseline algorithm.

```

1 Function Insert( $u, v$ ):
2    $B_2(u) \leftarrow B_2(u) \cup B_1(v)$ 
3   foreach  $w \in \mathcal{N}(u) \setminus \{v\}$  do
4      $B_2(w) \leftarrow B_2(w) \cup \{v\}$ 
5   end
6    $B_2(v) \leftarrow B_2(v) \cup B_1(u)$ 
7   foreach  $w \in \mathcal{N}(v) \setminus \{u\}$  do
8      $B_2(w) \leftarrow B_2(w) \cup \{u\}$ 
9   end
10 end

```

The magnitude of the changes (and the associated computational costs) induced by $\text{Insert}(u, v)$ vary. In particular, $B_2(u)$ can change significantly, as all vertices in $B_1(v)$ will be included in $B_2(u)$ (we refer to this as a *heavy* update, see Algorithm 1 at Line 2). Instead, for any vertex $w \in \mathcal{N}(u) \setminus \{v\}$, $B_2(w)$ will grow by at most one element, namely v (this is referred to as a *light* update, see Algorithm 1 at Line 4). Symmetrically, the same holds for v and for every $w \in \mathcal{N}(v) \setminus \{u\}$ (Lines 6 and 8).

A key observation at this point is that, while heavy updates can be addressed using (possibly approximate) data structures that allow efficient merging of 1- and 2-balls, this line of attack fails with light updates, whose cost derives from their potential number, which can be large in many real cases, as we noted in the introduction.

A first idea to reduce the average number of updates per edge insertion is to perform heavy updates immediately, instead processing light updates in batches that are performed occasionally. More precisely, when a new edge (u, v) arrives, it is initially marked as a *red* edge. Whenever the number of red edges incident to a vertex u exceeds a certain threshold, all the corresponding light updates are processed, and the state of red edges is updated to *black*. See Figure 1 for an example.




Figure 1: Example of insertion of a new edge (u, v) . The algorithm merges the 1-ball of v with the 2-ball of u (heavy update), while it does not immediately add vertex v to the 2-ball of vertex w or any other of u 's neighbors.

The idea behind the threshold-based approach is to maintain a balance between the number of black and red edges for every vertex. While useful when edge insertions appear in a random order, this approach may fail when red edges considerably expand the original size of the 2-ball of some vertex u . In order to mitigate this problem, our algorithm uses a second ingredient: upon each edge insertion (u, v) , the algorithm selects k vertices from $\mathcal{N}(u)$ and k from $\mathcal{N}(v)$ uniformly at random and performs a batch of light updates for the selected vertices, even if the threshold has not been reached yet.

These ideas are formalized in Algorithm 3. For each vertex v , our algorithm maintains two sets $\hat{B}_1(v)$ and $\hat{B}_2(v)$, as well as the *black degree* Δ_v and *red degree* δ_v of v . Our algorithm guarantees that $\hat{B}_1(v)$ is exactly $B_1(v)$, while $\hat{B}_2(v)$ is in general a subset of $B_2(v)$. The algorithm uses two global parameters, namely a *threshold* $\varphi \in [0, 1]$, and an integer k . The role of the parameter φ can be understood as follows: when φ is set to 0, the algorithm performs all heavy and light updates for every edge insertion, ensuring that $\hat{B}_2(v)$ always matches $B_2(v)$. As φ increases, the update function becomes lazier: light updates are not always executed, and $\hat{B}_2(v)$ is typically a proper subset of $B_2(v)$. For instance, when $\varphi = 1$, light updates are performed in batches every time the degree of a vertex doubles. Parameter k specifies the number of neighbors of v that are randomly selected for update of their 2-balls whenever an edge insertion involving v occurs. This mechanism corresponds to Lines 14 to 16 of Algorithm 3.

We call $\text{LAZY-ALG}(\varphi, k)$ the algorithm that runs Algorithm 2 on an initial graph $G^{(0)}$ and then processes a sequence S of edge insertions by running Algorithm 3 on each edge of S .

Algorithm 2: Init operation

Data: An undirected graph $G = (V, E)$, a threshold parameter $0 \leq \varphi \leq 1$, and an integer $k \geq 0$.

```

1 set  $\varphi$  and  $k$  as global parameters
2 foreach vertex  $u \in V$  do
3    $\delta_u \leftarrow 0$ 
4    $\Delta_u \leftarrow \deg_u$ 
5    $\hat{B}_1(u) \leftarrow B_1(u)$ 
6    $\hat{B}_2(u) \leftarrow B_2(u)$ 
7 end
```

A note on neighborhood representation. As we mentioned in the introduction, we treat $\hat{B}_1(v)$ and $\hat{B}_2(v)$ as sets of vertices in this section and in Section 4. We remark that this only serves the purpose of analyzing the error introduced by our lazy update policies: lossless representations of 1- and 2-balls may be unfeasible for medium or large graphs and compact data sketches are typically used to represent them in such cases. The choice of the actual sketch strongly depends on the type of query (or queries) one wants to support, such as 1- or 2-ball sizes [10, 28] or Jaccard similarity between 2-balls [6, 14, 21]. All sketches used for typical neighborhood queries are well-understood and come with strong accuracy guarantees. Moreover, they allow to perform the union of 1- and 2-balls we are interested in time proportional to the sketch size, which is independent of the sizes of the balls to merge [1].

2.2 Cost analysis for arbitrary sequences

Consistently to what we remarked above, our cost analysis focuses on the number of *set-union* operations: This performance measure in fact dominates the computational cost of Algorithm 3. More in detail, given any sequence S of edge insertions, starting from an initial graph $G^{(0)}$, we denote by $T(S)$ the overall number of union operations performed in Lines 4, 5, 11 and 16 of Algorithm 3 on the input sequence S .

Algorithm 3: INSERT

```

1 Function Insert( $(u, v)$ ):
2   for  $x \in \{u, v\}$  do
3     let  $y \in \{u, v\} \setminus \{x\}$ 
4      $\hat{B}_1(x) \leftarrow \hat{B}_1(x) \cup \{y\}$ 
5     // heavy update
6      $\hat{B}_2(x) \leftarrow \hat{B}_2(x) \cup \hat{B}_1(y)$ 
7      $\delta_x \leftarrow \delta_x + 1$ 
8     if  $\delta_x \geq \varphi \cdot \Delta_x$  then
9        $\Delta_x \leftarrow \Delta_x + \delta_x$ 
10       $\delta_x \leftarrow 0$ 
11      foreach  $z \in \mathcal{N}(x)$  do
12        // batch of light updates
13         $\hat{B}_2(z) \leftarrow \hat{B}_2(z) \cup \hat{B}_1(x)$ 
14      end
15    else
16      select  $k$  vertices  $w_1, \dots, w_k \in \mathcal{N}(x)$  u.a.r.
17      for  $i = 1, \dots, k$  do
18        // batch of light updates
19         $\hat{B}_2(w_i) \leftarrow \hat{B}_2(w_i) \cup \hat{B}_1(x)$ 
20      end
21    end
22  end
```

We observe that a trivial upper bound to $T(S)$ is $O(\Delta|S|)$, since each insertion can cost $O(\Delta)$ union operations where Δ is the maximum degree of the current graph. However, this trivial argument turns out to be too pessimistic: in what follows, we provide a more refined analysis of the amortized cost per edge insertion. We say that an algorithm has *amortized cost* \hat{c} per edge insertion if, for any sequence S of edge insertions, we have $T(S) \leq \hat{c}|S|$.

LEMMA 1. *Given any initial graph $G^{(0)}$ and any sequence S of edge insertions, the amortized update cost of Algorithm 3 is $O(\frac{1}{\varphi} + k)$ per edge insertion.*

PROOF. Let us first consider the case $k = 0$, i.e., when the random selection and the consequent instructions in Lines 14 to 16 are never performed. Our amortized analysis makes use of the *accounting method* [50]. The idea is paying the cost of any batch of light updates by charging it to previous edge insertions. More precisely, we assign *credits* to each edge insertion that we will use to pay the cost of subsequent batches of light updates. Formally, the *amortized cost* of an edge insertion is defined as the *actual cost* of the operation, plus the credits we assign to it, minus the credits (accumulated from previous operations) we spend for it. We need to carefully define such credits in order to guarantee that the sum of the amortized costs is an upper bound to the sum of the actual costs, i.e. we always have enough credits to pay for costly batch light updates.

We proceed as follows. When we insert the edge (u, v) , we put $2/\varphi$ credits on u and $2/\varphi$ credits on v . Now we bound the actual and amortized cost of each insertion.

First, consider an edge insertion (u, v) that does not trigger a batch of light updates. Its actual cost is 4 union operations (those in Lines 4 and 5, 2 for each endpoint of (u, v)). Then its amortized cost is upper-bounded by $4 + 4/\varphi = O(1/\varphi)$. Now consider the case in which the insertion causes a batch of light updates for u , or v , or both. We show that the credits accumulated by previous insertions are sufficient to pay for its cost. To see this, consider a batch of light updates involving vertex $x \in \{u, v\}$. And let Δ_x and δ_x be the current black and red degrees of x at that time (just before Line 7 is evaluated). It is clear that for vertex x we have accumulated $\delta_x \cdot 2/\varphi$ credits that now we use to pay for the cost of Lines 10 and 11. This cost equals to \deg_x union operations, where \deg_x is the current degree of vertex x . Since the batch of light updates has just been triggered, we have that $\delta_x \geq \varphi \Delta_x$, and hence we have at least $\delta_x \cdot 2/\varphi \geq \varphi \Delta_x \cdot 2/\varphi = 2\Delta_x$ credits to pay for the $\deg_x = \Delta_x + \delta_x = \Delta_x + \varphi \Delta_x \leq 2\Delta_x$ union operations. This concludes the proof.

Finally, to obtain the claim when $k > 0$, we notice that, in this case, every edge insertion causes $O(k)$ additional union operations. \square

3 RANDOM EDGE SEQUENCES

In this section, we analyze the accuracy of our lazy-update algorithm(s) over an arbitrary dynamic graph, whose edges are given in input as a uniformly sampled, random permutation over its edge set. Dynamic graphs resulting from random sequences of edge insertions have been an effective tool to provide theoretical insights that have often proved robust to empirical validation in various dynamic scenarios [18, 34, 41, 42, 44]. In more detail, assume $G = (V, E)$, with $|E| = t$, is the graph observed up to some time t of interest. Following [42, 44], we assume that the sequence of edges up to time t is chosen uniformly at random from the set of all permutations over E .¹⁰ The following fact is an immediate consequence of well-known and intuitive properties of random permutations. We state it informally for the sake of completeness, avoiding any further, unnecessary notation.

FACT 1. *Consider a dynamic graph $G = (V, E)$, whose edges are observed sequentially according to a permutation over E chosen uniformly at random. Then, for every $E' \subseteq E$, the sequence in which edges in E' are observed is itself a uniformly chosen, random permutation over E' .*

In the remainder, for an arbitrary vertex v , we analyze how well the output $\hat{B}_2(v)$ of Algorithm 3 approximates $B_2(v)$ at any round t in terms of its coverage:

DEFINITION 3.1. *We say that the output $\hat{B}_2(v)$ of $\text{LAZY-ALG}(\varphi, k)$ is a $(1 - \varepsilon)$ -covering of $B_2(v)$ if the following holds: i) $\hat{B}_2(v) \subseteq B_2(v)$; ii) $\mathbb{E} [|\hat{B}_2(v)|] \geq (1 - \varepsilon)|B_2(v)|$, where expectation is taken over the randomness of the algorithm and/or the input sequence. When the algorithm produces a $(1 - \varepsilon)$ -covering $\hat{B}_2(v)$ of $B_2(v)$ for every v , we say it has approximation ratio $\frac{1}{(1 - \varepsilon)}$.*

Our main result in this section is formalized in the following

¹⁰It should be noted that this includes the general case in which t is any intermediate point of a longer stream that possibly extends well beyond t . In this case, it is well-known and easy to see that, conditioned on the subset E of the edges released up to time t , their sequence is just a permutation of E .

THEOREM 1. *Let $\varepsilon \in (0, 1)$, and fix parameters $k = 0$ and $\varphi = \frac{\varepsilon}{1 - \varepsilon}$. Consider any graph $G(V, E)$ submitted as a uniform random permutation of its edge set E to $\text{LAZY-ALG}(\varphi, k)$. Then, at every time step $t \leq |E|$, the algorithm has approximation ratio $\frac{1}{1 - \varepsilon}$. Moreover, for every $\alpha > 0$ and every vertex $v \in V$, we have:*

$$\mathbb{P} \left(|\hat{B}_2(v)| < \frac{1 - \alpha}{1 + \varphi} |L_2(v)| \right) \leq e^{-\frac{2\alpha^2 |L_2(v)|}{(1 + \varphi)^2}}. \quad (1)$$

PROOF. Fix a vertex $v \in V$ and a round $t \geq 1$. In the remainder of this proof, all quantities are taken at time t . We are interested in how close $|\hat{B}_2(v)|$ is to $|B_2(v)|$. To begin, we note that the following relationship holds deterministically:

$$|\hat{B}_2(v)| = 1 + |L_1(v)| + |\hat{B}_2(v) \cap L_2(v)|, \quad (2)$$

where the only random variable on the right hand side is the last term. We next define a partition $C = \{C_u : u \in L_1(v)\}$ of $L_2(v)$ as




Figure 2: Example of a partition of $L_2(v)$ into three sets C_1, C_2, C_3 . Edges connecting vertices $w \in L_2(v)$ to their respective partitions are thicker.

follows: for each $w \in L_2(v)$, we choose a vertex $u \in L_1(v) \cap \mathcal{N}(w)$ and assign w to C_u . This way, each vertex $w \in L_2(v)$ is associated to exactly *one* edge connecting one vertex in $L_1(v)$ to w (see Figure 2, where the edges in question are thick in the picture). Let E_v denote the set of such edges and note that i) E_v is a subset of the edges connecting vertices in $L_1(v)$ to those in $L_2(v)$, ii) $|E_v| = |L_2(v)|$ by definition and iii) $|C_u| \leq \deg_u - 1$ for every $u \in L_1(v)$, given that C_u contains a subset of u 's neighbors and (v, u) is always present. Moreover, Algorithm 3 guarantees that $|\hat{B}_2(v) \cap L_2(v)|$ is at least the number of edges in E_v that are black. These considerations allow us to conclude that

$$|\hat{B}_2(v) \cap L_2(v)| \geq |\{e \in E_v : e \text{ is black}\}|.$$

A key observation at this point is that Algorithm 3 implies that for every $x \in V$, $\delta_x \leq \left\lfloor \frac{\varphi}{1 + \varphi} \deg_x \right\rfloor$. As a consequence, if some $e = (u, w) \in E_v$ was not among the last $\left\lfloor \frac{\varphi}{1 + \varphi} \deg_u \right\rfloor$ edges incident in u that were released within time t , it is necessarily black. For $e = (u, w) \in E_v$, with $u \in L_1(v)$ and $w \in L_2(v)$, let $X_e = 1$ if e was among the first $\deg_u - \left\lfloor \frac{\varphi}{1 + \varphi} \deg_u \right\rfloor$ edges incident in u that were released up to time t and let $X_e = 0$ otherwise. Following the argument above, the event $(X_e = 1)$ implies the event " e is black", whence:

$$|\hat{B}_2(v) \cap L_2(v)| \geq |\{e \in E_v : e \text{ is black}\}| \geq \sum_{e \in E_v} X_e. \quad (3)$$

Next, we are interested in bounds on $\mathbb{P}(X_e = 1)$. Assume e is incident in u and let S be the set of edges incident in u observed up

to time t . Then, from Fact 1, the sequence in which these edges are observed is just a random permutation of S . This immediately implies that, if e is incident to a vertex $u \in L_1(v)$, then

$$\mathbb{P}(X_e = 1) = \frac{\deg_u - \left\lfloor \frac{\varphi}{1+\varphi} \deg_u \right\rfloor}{\deg_u} \geq \frac{1}{1+\varphi}.$$

Together with (3) this yields:

$$\mathbb{E}[|\hat{B}_2(v)|] \geq 1 + |L_1(v)| + \frac{1}{1+\varphi} |L_2(v)| \geq \frac{1}{1+\varphi} |B_2(v)|.$$

We next show that $\sum_{e \in E_v} X_e$ is concentrated around its expectation when $|L_2(v)|$ is large enough, which implies that $|\hat{B}_2(v)|$ is concentrated around a value close to $|B_2(v)|$ in this case. The main technical hurdle here is that the X_e 's are correlated (albeit mildly, as we shall see). To prove concentration, we resort to Martingale properties of random edge sequences to apply the method of (Average) Bounded Differences [23]. In order to do this, we need bounds on $\mathbb{P}(X_e = 1 | X_f = 1)$ and $\mathbb{P}(X_e = 1 | X_f = 0)$, for $e, f \in E_v$. Assume again that e is incident in $u \in L_1(v)$ and that S is the set of edges incident in u observed up to time t . Assume first that f is also incident in u and that, without loss of generality, f is the i -th edge to appear among those in S . $X_f = 1$ implies $i \leq \deg_u - \left\lfloor \frac{\varphi}{1+\varphi} \deg_u \right\rfloor$. On the other hand, for any such choice for f 's position in the sequence, Fact 1 implies that e will appear in a position j that is sampled uniformly at random from the remaining ones, so that $\mathbb{P}(X_e = 1 | X_f = 1) = \frac{\deg_u - \left\lfloor \frac{\varphi}{1+\varphi} \deg_u \right\rfloor - 1}{\deg_u - 1}$ in this case. With a similar argument, it can be seen that $\mathbb{P}(X_e = 1 | X_f = 0) = \frac{\deg_u - \left\lfloor \frac{\varphi}{1+\varphi} \deg_u \right\rfloor}{\deg_u - 1}$. Intuitively and unsurprisingly, the events $(X_e = 1)$ and $(X_f = 1)$ are negatively correlated, while $(X_e = 1)$ and $(X_f = 0)$ are positively correlated. This allows us to conclude that $\mathbb{P}(X_e = 1 | X_f = 1) \leq \mathbb{P}(X_e = 1 | X_f = 0)$ and

$$\mathbb{P}(X_e = 1 | X_f = 0) - \mathbb{P}(X_e = 1 | X_f = 1) \leq \frac{1}{\deg_u - 1}.$$

Assume next that f is not incident in u . Again from Fact 1, in this case f has no bearing on the relative order in which edges incident in u appear, so that $\mathbb{P}(X_e = 1 | X_f = 0) = \mathbb{P}(X_e = 1 | X_f = 1) = \mathbb{P}(X_e = 1)$. Now, without loss of generality, suppose that $f = (z, w)$, with $z \in L_1(v)$, so that $w \in C_z$. Denote by $E_v(z)$ the subset of edges in E_v with one end point in C_z . Moving to conditional expectations

we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{e \in E_v} X_e \mid X_f = 0 \right] - \mathbb{E} \left[\sum_{e \in E_v} X_e \mid X_f = 1 \right] \\ &= \sum_{e \in E_v} \left(\mathbb{P}(X_e = 1 \mid X_f = 0) - \mathbb{P}(X_e = 1 \mid X_f = 1) \right) \\ &= \sum_{e \in E_v \setminus E_v(z)} \left(\mathbb{P}(X_e = 1 \mid X_f = 0) - \mathbb{P}(X_e = 1 \mid X_f = 1) \right) \\ &+ \sum_{e \in E_v(z)} \left(\mathbb{P}(X_e = 1 \mid X_f = 0) - \mathbb{P}(X_e = 1 \mid X_f = 1) \right) \\ &\leq \frac{|C_z|}{\deg_z - 1} \leq 1, \end{aligned}$$

where the third inequality follows from the definition of C_z , since f is incident in z , while the last inequality follows since $|C_z| \leq \deg_z - 1$ for every $z \in L_1(v)$, because one of the edges incident in z is by definition the one shared with v .

We can therefore apply [23, Definition 5.5 and Corollary 5.1], with $c \leq |L_2(v)|$ to obtain, for every $\alpha > 0$:

$$\mathbb{P} \left(\mathbb{E} \left[\sum_{e \in E_v} X_e \right] - \sum_{e \in E_v} X_e > \alpha \mathbb{E} \left[\sum_{e \in E_v} X_e \right] \right) \leq e^{-\frac{2\alpha^2 |L_2(v)|}{(1+\varphi)^2}},$$

where in the right hand side we also used the bound $\mathbb{E}[\sum_{e \in E_v} X_e] \geq \frac{1}{1+\varphi} |L_2(v)|$ we showed earlier. Finally, we recall (2) and (3) to conclude that $|\hat{B}_2(v)| \geq \frac{1-\alpha}{1+\varphi} |L_2(v)|$ with (at least) the same probability.

Theorem 1 easily implies approximation bounds on indices that depend on the union and/or intersection of 2-balls. For example, we immediately have the following approximation bound on the Jaccard similarity between any pair of 2-balls.

COROLLARY 2. *Under the same assumptions as Theorem 1, at any time step $t \geq 1$ and for any pair of vertices $u, v \in V$, LAZY-ALG(φ, k) guarantees the following approximation of the Jaccard similarity*

$$\begin{aligned} & \text{between } B_2(u) \text{ and } B_2(v) \text{ with probability at least } 1 - e^{-\frac{2\alpha^2 |L_2(u)|}{(1+\varphi)^2}} - \\ & e^{-\frac{2\alpha^2 |L_2(v)|}{(1+\varphi)^2}} : \\ & \frac{J(B_2(u), B_2(v))}{1 - 2\epsilon'} \geq J(\hat{B}_2(u), \hat{B}_2(v)) \geq (1 - \epsilon') J(B_2(u), B_2(v)) - \epsilon', \end{aligned} \tag{4}$$

where $\epsilon' = \frac{\varphi + \alpha}{1 + \varphi}$.

PROOF. It is easy to see that $|\hat{B}_2(u)| \geq (1 - \epsilon') |B_2(u)|$ and $|\hat{B}_2(v)| \geq (1 - \epsilon') |B_2(v)|$ together imply (4) deterministically. The result then immediately follows from Theorem 1 and a union bound on the events $(|\hat{B}_2(u)| < (1 - \epsilon') |L_2(u)|)$ and $(|\hat{B}_2(v)| < (1 - \epsilon') |L_2(v)|)$. \square

4 ADVERSARIAL EDGE SEQUENCES

We next study our lazy-update algorithm in an adversarial framework. We show in Section 4.1 that if the adversary can *both*: i) choose a worst-case graph G and ii) submit G according to an *adaptive* sequence of edge insertions, then it is possible to prove a

strong lower bound on the achievable update-time/approximation trade-off of the whole parameterized scheme $\text{LAZY-ALG}(\varphi, k)$.

On the other hand, in Section 4.2 we provide a *necessary* condition for the adversarial, worst-case framework above: the *girth* [22] of G must be smaller than 5. More precisely, G must contain an unbounded number of triangles or cycles of length 4. We do this by showing that for a suitable parameter setting, algorithm $\text{LAZY-ALG}(\varphi, k)$ achieves almost-optimal trade-offs even on adversarial edge insertion sequences, for every graph that has a bounded number of such small cycles (see Definition 4.1 for a formal definition of this class of graphs).

4.1 A lower bound for adversarial sequences

The lower bound for the adversarial framework described above is formalized in the following result on the approximation ratio (see Def. 3.1)

THEOREM 3. *For every $\varphi \in [0, 1]$, and integer $k \geq 0$, if $\text{LAZY-ALG}(\varphi, k)$ has approximation ratio $\rho \geq 1$, then it must have an amortized update cost of $\Omega(\Delta/\rho^3)$, where Δ is the maximum degree of the graph.*




Figure 3: Black edges are present at $t = 0$, while red ones are inserted in the interval $\{1, 2, \dots, \Delta\rho^2\}$. At time $t > 0$, an edge with one endpoint in $u_{1t \bmod \Delta}$ and the other in a distinct 0-degree vertex in S_2 is added.

PROOF. Fix $\rho \geq 1$, and assume that $\text{LAZY-ALG}(\varphi, k)$ has an approximation ratio of at most ρ . We will show that there exist an initial graph $G^{(0)}$ with degree Δ and a sequence of edge insertions against which $\text{LAZY-ALG}(\varphi, k)$ must incur an amortized update time of $\Omega(\Delta/\rho^3)$.

Note that for $k > 0$, the algorithm is randomized. In order to address this, we prove our lower bound for every possible realization of the randomness used by the algorithm. Therefore, we assume the values of the random bits used by $\text{LAZY-ALG}(\varphi, k)$ are fixed arbitrarily (and optimally) and we assume henceforth that the behavior of the algorithm is completely deterministic.

The initial graph $G^{(0)}$ consists on 2Δ vertices forming a complete bipartite graph with sides $S_0 = \{u_{01}, \dots, u_{0\Delta}\}$ and $S_1 = \{u_{11}, \dots, u_{1\Delta}\}$, along with an additional set S_2 of $\Delta\rho^2$ isolated vertices (see Figure 3). The sequence of edge insertions is defined as follows: for each vertex in S_1 , we insert ρ^2 new edges. Each of these $\Delta\rho^2$ edges connects a vertex in S_1 to a previously isolated vertex in S_2 .

Consider the time instant right after all edge insertions. Since we assumed that the algorithm guarantees an approximation ratio

of ρ , it holds that for every vertex $u \in S_0$, $|\hat{B}_2(u)| \geq \frac{1}{\rho} |B_2(u)| = \frac{1}{\rho} (2\Delta + \Delta\rho^2) = \Delta\rho + 2\Delta/\rho$. This implies that after all edge insertions, u must be aware of at least $\Delta\rho + 2\Delta/\rho - 2\Delta = \Delta\rho - 2\Delta(1 - 1/\rho)$ vertices from S_2 .

We say that there is a *message* from v to u if vertex v performs a union operation of the form $\hat{B}_2(u) \leftarrow \hat{B}_2(u) \cup \hat{B}_1(v)$.

Since, at any time, every vertex $v \in S_1$ is adjacent to at most ρ^2 vertices in S_2 , it must be that each $u \in S_0$ must have received at least $\frac{\Delta\rho - 2\Delta(1 - 1/\rho)}{\rho^2} = \Omega(\Delta/\rho)$ messages from vertices in S_1 . As a consequence, the total number of messages are at least $\Omega(\Delta^2/\rho)$. As the number of insertions is $\Delta\rho^2$, the amortized update cost per insertion is at least $\frac{\Omega(\Delta^2/\rho)}{\Delta\rho^2} = \Omega(\Delta/\rho^3)$. \square

We have special cases as corollaries. We need amortized update cost $\Omega(\Delta)$ if we want $\rho = O(1)$, $\Omega(\sqrt[4]{\Delta})$ if we want $\rho = O(\sqrt[4]{\Delta})$ and so on.

REMARK 1. *The lower bound in Section 4.1 in fact holds for a wider class of algorithms. Informally speaking, this class includes any local algorithm that limits its online updates to the 2-hop neighbors of u and v only. Making this claim more formal requires addressing several technical issues that are outside the scope of the present work.*

4.2 Locally γ -sparse graphs

In this section, we provide the characterization of a class of graphs for which our lazy-update approach always guarantees good amortized cost/approximation trade-offs, even under the assumption of adversarial edge insertion sequences. Given a graph $G(V, E)$ and a subset $V' \subseteq V$, we denote by $G[V']$ the subgraph induced by V' . Informally, a graph is locally γ -sparse if every node in $B_2(u) \setminus \{u\}$ has roughly at most γ neighbors in $L_1(u)$.

DEFINITION 4.1 (LOCALLY γ -SPARSE GRAPHS). *Let $\gamma \in \{0, 1, \dots, n-2\}$. A graph $G(V, E)$ is said locally γ -sparse if for each vertex $u \in V$:* (i) $\forall v \in L_1(u)$ the degree of v in $G[L_1(u)]$ is at most γ , and (ii) $\forall w \in L_2(u)$ the degree of w in $G[L_1(v) \cup \{w\}]$ is at most $\gamma + 1$.

Observe that the class of locally γ -sparse graphs grows monotonically with γ , including all possible graphs for $\gamma = n - 2$, while the most restricted class is obtained for $\gamma = 0$. It is interesting to note that locally γ -sparse graphs are not necessarily sparse in absolute terms. For example, for $\gamma = 0$, the class coincides with the well-known class of graphs with *girth* at least 5: these graphs can have up to $\Theta(n^{\frac{3}{2}})$ edges assuming Erdős' Girth Conjecture [25] (the proof of such equivalence is given in the full version [8]).

A first, preliminary analysis of our lazy-update approach considers the deterministic version of Algorithm 3, i.e., when $k = 0$. It turns out that this version achieves an approximation ratio of $\frac{\gamma+1}{1-\epsilon}$ and amortized cost $O(1/\epsilon)$ (see full version [8]). So, the approximation accuracy decreases linearly in the parameter γ . Interestingly enough, we instead show that a suitable number of random light updates allows Algorithm 3 to perform much better than its deterministic version. This is the main result of this section and it is formalized in the next theorem, whose proof is available in the full version of this paper [8].

THEOREM 4. *Let $\epsilon \in (0, 1)$, and let $G^{(0)}$ be an initial graph. Consider any sequence of edge insertions that yields a final graph G . If G*

is locally γ -sparse $\text{LAZY-ALG}(\varphi = 1, nk = \frac{4(\gamma+1)}{\epsilon})$ has approximation ratio of $\frac{1}{1-\epsilon}$ and amortized cost $O\left(\frac{\gamma+1}{\epsilon}\right)$.

5 EXPERIMENTAL ANALYSIS

The purpose of this experimental analysis, is the validation of our lazy approach along three main axes: i) the consistency between the theoretical findings from Sections 2 to 4 and the actual behavior of our lazy approach on real datasets; ii) the accuracy of our approach on two key neighborhood-based queries, namely, size and Jaccard similarity; iii) its computational savings with respect to non-lazy baselines on medium and large incremental graphs. For the sake of space, the complete set of results appears in the full version [8].

5.1 Experimental setup

Platform. Our experiments were performed on a machine with 2.3 GHz Intel Xeon Gold 5118 CPU with 24 cores, 192 GB of RAM, cache L1 32KB, shared L3 of 16MB and UMA architecture. The whole code is written in C++, compiled with GCC 10 and with the following compilation flags: `-DARCH_X86_64 -Wall -Wextra -pg -O3 -lm`.¹¹

Algorithms. We compared $\text{LAZY-ALG}(\varphi, k)$ with various combinations (φ, k) with the following baselines:

- i) the exact *algorithm 1*, which is not scalable and is only used in the first round of experiments on smaller datasets, in order to isolate the error introduced by lazy updates;
- ii) the naive *sketch-based baseline*, which adopts sketch-based representations of 1- and 2-balls, but performs all necessary updates. It corresponds to Algorithm 3 with $\varphi = 0$ and $k = 0$, but where 1- and/or 2-ball unions correspond to merging the corresponding sketches.¹²

With the exception of the first set of experiments (i.e. item i) above), we compared our algorithm to the sketch-based baseline. For some results, we needed to compute the true value of the parameter of interest for 2-balls (e.g., size) by executing a suitably optimized BFS. For our LAZY-ALG , we used combinations of the following values: $\varphi = 0.1, 0.25, 0.5, 0.75, 1$ and $k = 0, 2, 4, 8$.

Implementation details. To best assess the performance of algorithms, it would be ideal to minimize the overhead deriving from the management of edge insertions in the graph. We observe that this overhead is the same for all algorithms we tested. Hence, we represented graphs using the *compressed sparse row* format [24] and, since we knew the edge insertion sequence in advance, we pre-allocated the memory needed to accommodate them, so as to minimize overhead. For experiments that required hash functions, we used *tabulation hashing* [45].

Datasets. We considered real incremental graphs of different sizes, both directed and undirected, available from NetworkRepository [48]. In our time analysis, we also extracted a large incremental dataset from a large static graph, namely soc-friendster [54], available from SNAP [39]. Following previous work on dynamic graphs [19, 32], we generated an incremental graph by adding edges

sequentially and in random order, starting from an empty graph. The main features of these datasets are summarized in Table 1.

5.2 Results

We focused on the case of undirected graphs, but our approach extends seamlessly to directed graphs, such as some of the real examples we consider in this section. In such cases, the only caveat to keep in mind is that we define the h -ball of a vertex u as the subset of vertices that are reachable from u over a directed path traversing at most h edges.

Impact of lazy updates. The goal of our first experiment is twofold: i) assessing the impact of lazy updates on the estimation of 2-balls; ii) assessing the degree of consistency between the theoretical findings of Sections 3 and 4 and the actual behavior of our algorithms on real datasets. In order to isolate the specific contribution of lazy updates in the estimation error, we implemented (true and approximate) 1-balls and 2-balls losslessly, as dictionaries. This way, the error in 2-ball size estimation is only determined by our lazy update policy. Since, as we argued elsewhere in the paper, lossless representations of 2-balls quickly becomes unfeasible for larger datasets, this first experiment was run on 3 small-medium datasets, namely, comm-linux-kernel-reply, fb-wosn-friends and ia-enron-email-all (results for comm-linux-kernel-reply are reported in the full version [8]).

For each of the above graphs, we selected as a sample 5000 vertices whose 2-balls are the largest at the end of the edge insertion sequence. For every pair (φ, k) of parameter values for $\text{LAZY-ALG}(\varphi, k)$, we performed 10 independent runs. Each run is organized into the following steps: 1) the initial graph $G^{(0)}(V, E^{(0)})$ corresponds to the first 20% edge insertions; 2) we measure the coverage (see Definition 3.1) of each of the 5000 2-balls above by $\text{LAZY-ALG}(\varphi, k)$ at each of the 100 equally spaced timestamps, the same in each run. For the generic timestamp t , we measure the average coverage

$$C_t = \frac{1}{5000} \sum_{i=1}^{5000} \frac{\hat{B}_i}{B_i},$$

where B_i and \hat{B}_i respectively denote the true and estimated sizes of the i -th 2-ball from the sample. Finally, for each timestamp t , we plot the average of the 10 values of C_t computed in every run. The results, summarized in Figure 4, are fully consistent with our theoretical findings from Section 3. At least for the diverse dataset sample considered here, uniform random permutations are a reasonable theoretical proxy of real sequences. More in general, real sequences seem to be rather far from the pathological worst-cases analyzed in Section 4.1, so that the actual behavior of our algorithm is not only in line, but better than our analysis predicts. We also have an initial insight into the effects of the parameters φ and k , which will be examined more thoroughly in the subsequent subsections.

Ball size estimation via sketches. The goal of the next round of experiments was assessing the accuracy of our algorithms in 2-ball size estimation, in the realistic setting in which approximate 1- and 2-balls are represented via state-of-art sketches that are based on probabilistic counters [5]. In this case, we have a compound estimation

¹¹Code available at https://github.com/Gnumlab/graph_ball.

¹²Again, the specific sketch (or sketches) used depends on the neighborhood queries one wants to support.

Table 1: Summary table of real networks used in the experiments.

Dataset	$ V $	$ E $	Directed	Dynamic
comm-linux-kernel-reply	27,927	242,976	✓	✓
fb-wosn-friends [51]	63,731	817,090	✗	✓
ia-enron-email-all	87,273	321,918	✓	✓
soc-flickr-growth	2,302,925	33,140,017	✓	✓
soc-youtube-growth	3,223,589	9,376,594	✓	✓
soc-friendster [54]	65,608,366	1,806,067,135	✗	✗

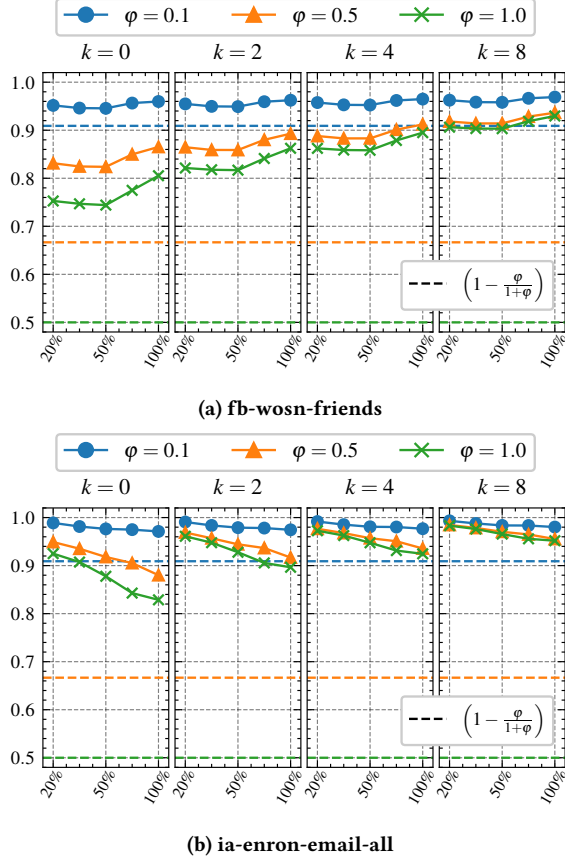


Figure 4: Average coverage C_t from a network with 20% of its edges to the end of the insertion sequence. Dashed lines show the theoretical expected coverage for random sequences (see Theorem 1).

error, arising from both our lazy update policy and the use of probabilistic counters. We ran experiments on the three small-medium sized datasets considered previously, plus two large ones, namely, soc-youtube-growth and soc-flickr-growth. Again, we considered the top-5000 largest 2-balls as sample. The experiments were executed as in the previous case, with the following differences: i) we considered 3 timestamps, respectively corresponding to 50%, 75% and 100% of all edge insertions of each dataset; ii) in this case, 1- and 2-balls are not explicitly represented as sets (not even by

the baseline). As a result, ball sizes might be overestimated and coverage has no clear meaning. We therefore use Mean Absolute Percentage Error (MAPE) to measure accuracy, defined as follows:

$$\text{MAPE} = \frac{1}{5000} \sum_{i=1}^{5000} \frac{|B_i - \hat{B}_i|}{B_i},$$

where i refers to the 2-ball size of the i -th sampled vertex. For each dataset, this index is computed for each of 10 independent runs at each of the 3 timestamps we consider. Results are summarized for all datasets and combinations (φ, k) we consider in Table 2 for the last timestamp (100% of edge insertions). Results for other timestamps (50% and 75%) are similar and are omitted for the sake of space (see full version [8] for complete results). The main takeaways here are that i) the additional error introduced by the use of sketches (which can be controlled by varying the size of the sketch) is relatively modest; ii) even relatively large values of φ and/or small values of k result in performances that are close to those of the baseline that uses sketches to represent balls, but naively performs all light updates. The effect of k is more pronounced when φ is large, contributing to a reduction in both error and, to a lesser extent, variance. Finally, we note that although the effect of increasing k can be achieved by decreasing φ , our analysis in Section 4 suggests that the parameter k provides robustness against worst-case scenarios. Even though the analysis of the impact of sketches on error is out of the scope of this paper, the results in the previous paragraph and the small difference with the baseline (Table 2) when using the sketches further suggest that the analysis of random sequences may also apply to real sequences.

Accuracy in Jaccard similarity estimation. With the same goal as in the previous experiment, we now evaluate the quality of our lazy approach policy combined with the use of sketches in another graph mining task: the Jaccard similarity estimation for 2-hop neighborhoods. The sketch used to represent the 1- and 2- balls is the well-known h -minhash [13], with $h = 100$ hash functions.

As before, 10 independent runs were performed for each combination of parameters φ, k , and the baseline, on each of the previously used datasets. Errors were measured at 3 different timestamps, corresponding to 50%, 75%, and 100% of the edge insertion sequence.

Considering all vertex pairs in the entire graph would be computationally prohibitive, as their number is too large. Moreover, many of these pairs would have an extremely low Jaccard similarity, making it difficult to estimate them with a reasonably small error [14, 15]. Therefore, to better evaluate our algorithm's and baseline's quality, we need vertex pairs with a sufficiently high similarity in

Table 2: Mean and standard deviation of absolute percentage errors for 2-hop neighborhood size estimation. Size estimates were made using the KMV probabilistic counter [5], with size 32. Queries were made at the end of the insertion sequence.

	k	$\varphi = 0.1$	$\varphi = 0.5$	$\varphi = 1$	baseline
linux	0	0.14 ± 0.12	0.19 ± 0.14	0.17 ± 0.11	0.12 ± 0.10
	2	0.13 ± 0.11	0.14 ± 0.10	0.16 ± 0.11	
	4	0.14 ± 0.09	0.17 ± 0.12	0.14 ± 0.10	
	8	0.14 ± 0.11	0.14 ± 0.09	0.13 ± 0.10	
fb-wosn	0	0.16 ± 0.12	0.15 ± 0.11	0.21 ± 0.12	0.14 ± 0.11
	2	0.13 ± 0.09	0.15 ± 0.10	0.17 ± 0.11	
	4	0.14 ± 0.11	0.15 ± 0.10	0.16 ± 0.11	
	8	0.16 ± 0.13	0.14 ± 0.10	0.15 ± 0.11	
enron	0	0.13 ± 0.10	0.16 ± 0.11	0.20 ± 0.14	0.13 ± 0.11
	2	0.14 ± 0.11	0.16 ± 0.12	0.16 ± 0.12	
	4	0.13 ± 0.12	0.15 ± 0.12	0.15 ± 0.12	
	8	0.13 ± 0.11	0.14 ± 0.10	0.16 ± 0.13	
flickr	0	0.17 ± 0.14	0.18 ± 0.12	0.17 ± 0.11	0.17 ± 0.14
	2	0.16 ± 0.12	0.12 ± 0.09	0.14 ± 0.10	
	4	0.14 ± 0.09	0.13 ± 0.10	0.16 ± 0.10	
	8	0.14 ± 0.11	0.13 ± 0.10	0.14 ± 0.09	
youtube	0	0.15 ± 0.11	0.15 ± 0.10	0.24 ± 0.11	0.14 ± 0.11
	2	0.16 ± 0.13	0.14 ± 0.10	0.19 ± 0.11	
	4	0.13 ± 0.11	0.13 ± 0.10	0.15 ± 0.11	
	8	0.13 ± 0.10	0.12 ± 0.09	0.15 ± 0.11	

their 2-hop neighborhoods. To address this, we adopted a similar sampling process as before: we selected the 5000 vertices with the largest 2-hop neighborhoods at the end of the edge insertion sequence and randomly chose 1000 pairs whose similarity is at least 0.2. On that sample, we evaluated the MAPE of the Jaccard similarity estimate computed using the h -minhash signatures. Table 3 reports the results at the end of the insertion sequence, while results for the other timestamps (50% and 75%) are similar and are omitted for the sake of space (see full version [8] for complete results).

This experiment further confirms the observations from the previous one: when using sketches to represent the 2-balls, the errors obtained with our lazy update policy (with appropriate choices of parameters φ, k) are similar and fully comparable to those of the baseline, which performs all necessary updates.

Run time analysis. Finally, we measured the running times of our algorithm and the corresponding speed-up with respect to the baseline. In the following, we report and discuss the results for the task of size estimation, using probabilistic counters. The performances for the task of Jaccard similarity estimation are analogous, and thus reported in the full version [8].

Table 4 reports the average speed-up of our algorithm with respect to the naive baseline for the same combinations (φ, k) considered previously, for all datasets except com-friendster. Speed-ups are computed in terms of *total update time*, i.e., the total time it takes to process the whole insertion sequence. Table 5 reports the overall processing times on com-friendster, for a subset of the

Table 3: Mean and standard deviation of absolute percentage errors for Jaccard similarity estimation, with 100 hash functions. Queries were made at the end of the insertion sequence.

	k	$\varphi = 0.1$	$\varphi = 0.5$	$\varphi = 1$	baseline
linux	0	0.11 ± 0.09	0.13 ± 0.09	0.11 ± 0.09	0.10 ± 0.08
	2	0.09 ± 0.07	0.12 ± 0.09	0.11 ± 0.09	
	4	0.10 ± 0.08	0.12 ± 0.09	0.10 ± 0.08	
	8	0.09 ± 0.07	0.12 ± 0.09	0.10 ± 0.08	
fb-wosn	0	0.70 ± 0.25	0.72 ± 0.24	0.74 ± 0.23	0.70 ± 0.24
	2	0.70 ± 0.24	0.70 ± 0.25	0.72 ± 0.24	
	4	0.70 ± 0.24	0.70 ± 0.24	0.72 ± 0.24	
	8	0.70 ± 0.24	0.70 ± 0.24	0.71 ± 0.25	
enron	0	0.09 ± 0.09	0.14 ± 0.12	0.15 ± 0.12	0.09 ± 0.09
	2	0.10 ± 0.09	0.12 ± 0.10	0.14 ± 0.12	
	4	0.09 ± 0.09	0.11 ± 0.10	0.13 ± 0.11	
	8	0.10 ± 0.09	0.10 ± 0.09	0.11 ± 0.10	
flickr	0	0.12 ± 0.09	0.11 ± 0.09	0.11 ± 0.09	0.11 ± 0.09
	2	0.11 ± 0.08	0.10 ± 0.08	0.11 ± 0.09	
	4	0.11 ± 0.08	0.10 ± 0.08	0.10 ± 0.08	
	8	0.11 ± 0.08	0.10 ± 0.08	0.11 ± 0.08	
youtube	0	0.12 ± 0.09	0.11 ± 0.09	0.17 ± 0.13	0.11 ± 0.09
	2	0.11 ± 0.09	0.11 ± 0.09	0.15 ± 0.11	
	4	0.11 ± 0.09	0.13 ± 0.10	0.16 ± 0.11	
	8	0.11 ± 0.09	0.12 ± 0.09	0.14 ± 0.10	

combinations of φ and k . It should be noted that the baseline did not complete within a reasonable amount of time in this case. Finally, Figure 5 illustrates the average time cost per operation for the baseline, as well as the slowest and fastest parameter settings of φ and k in our algorithm. These experiments clearly highlight that the number of lazy updates is the crucial factor affecting performance and that our approach is very effective at addressing this problem, resulting in considerable speed-ups that grow with the size of the graph. These results regarding processing times, in conjunction with previous findings, underscore the significant advantage in terms of speed at the expense of a modest and acceptable reduction in query quality which, due to the necessary use of sketches, is never totally accurate.

A quick reference for parameter setting. The trade-off between accuracy, parameter setting and update costs can be summarized as follows: decreasing φ increases the frequency of heavy updates (and thus accuracy), while increasing k increases the number of light updates upon a single edge insertion, again improving accuracy. As our experimental results show, real edge sequences exhibit a behavior that is largely consistent with the random permutation model adopted in the analysis proposed in Section 3 and far from the adversarial sequences analyzed in Section 4. In this perspective, the choice of the parameters can be derived by our theoretical results (see Theorem 1 and Lemma 1). However, in practice, by running our algorithms on a diverse set of real networks for a multiplicity of parameter settings, we identified two robust settings, corresponding




Figure 5: Average time per insertion operation (in seconds) in log-scale. The datasets are arranged on the x-axis in increasing order of number of operations, again in log-scale. The baseline time for soc-friendster dataset is not reported, since it exceeded a time limit of 36 hrs.

Table 4: Speed up with respect to the baseline, using KMV probabilistic counters [5], with size 32.

	k	φ				
		0.1	0.25	0.5	0.75	1
linux	0	14.39x	22.73x	30.27x	32.36x	35.65x
	2	12.00x	18.22x	22.03x	21.49x	23.18x
	4	11.52x	15.34x	17.63x	16.93x	17.54x
	8	9.57x	11.86x	12.80x	12.03x	12.22x
fb-wosn	0	5.48x	9.40x	11.15x	15.23x	16.22x
	2	5.22x	7.35x	7.16x	9.21x	9.67x
	4	4.32x	6.14x	6.05x	7.18x	7.23x
	8	3.79x	4.38x	4.91x	5.31x	5.22x
enron	0	4.17x	5.38x	6.64x	7.12x	7.70x
	2	3.61x	4.11x	4.68x	4.82x	5.27x
	4	3.05x	3.42x	3.78x	3.96x	4.21x
	8	2.51x	2.73x	2.85x	3.01x	3.12x
flickr	0	13.52x	19.97x	26.54x	31.51x	34.20x
	2	12.16x	16.41x	19.52x	21.93x	22.61x
	4	10.76x	13.96x	16.35x	17.17x	17.79x
	8	9.10x	11.19x	12.36x	12.99x	13.28x
youtube	0	41.61x	61.80x	77.08x	86.50x	93.73x
	2	38.16x	51.64x	60.39x	64.93x	65.51x
	4	36.08x	47.43x	52.26x	55.06x	55.90x
	8	33.33x	39.60x	42.73x	43.61x	43.82x

to different choices for the trade-off between accuracy and efficiency. The first setting is $\varphi = 0.5$ and $k = 0$, which yields a very fast solution with a reasonable accuracy. Instead, when the accuracy requirement is more stringent, we suggest setting $\varphi = 0.25$ and $k = 2$, which usually provides a negligible error with a reasonable speed-up.

Finally, the setting of the parameters for the sketches we used (in particular, their sizes) was simply driven by i) the trade-off established by the theoretical analyses in [5] and [13, 14, 20] between

Table 5: Total running time for ball size estimation using probabilistic counters, for com-friendster dataset. The time for the baseline algorithm is not reported since it exceeded a time limit of 36 hrs.

k	φ				
	0.1	0.25	0.5	0.75	1
0	5h 2'	2h 55'	1h 53'	1h 27'	1h 18'
2	6h 3'	3h 48'	2h 59'	2h 42'	2h 34'
4	6h 50'	4h 42'	3h 56'	3h 39'	3h 32'
8	8h 23'	6h 53'	5h 54'	5h 30'	5h 19'

accuracy and sketch sizes and ii) the amount of memory we could afford. Ultimately, this resulted in the choices described in Section 5.1, where we used 32-bit integers.

6 OUTLOOK

This work leaves a number of open questions that might deserve further investigation. A first, obvious direction is extending our approach to handle edge deletions. Actually, a simple strategy to manage a deletion of an edge (u, v) is to recompute the 2-balls of the $O(\deg_u + \deg_v)$ affected vertices via 2-layer BFS visits. The cost of this straightforward approach can be amortized on the entire sequence of edge updates whenever deletions are rare. However, the general case is much more challenging: while our approach is potentially useful, one of the main problems here is handling deletions when compact, sketch-based data structures are used to represent 1- and 2-balls. Some recent contributions address similar issues in dynamic data streams [17, 20], but extending our analyses to this general case does not seem straightforward. Another interesting direction is investigating strategies to handle queries over h -balls when $h > 2$, for example maintaining their sizes under dynamic updates. In this case, each edge addition/deletion potentially has cascading effects over h -hops. Optimizing (amortized) update costs in this general setting does not seem trivial and we conjecture that a dependence on h might be necessary. Finally, we remark that our algorithms are inherently local, i.e., whether or not to perform updates involving any vertex v only depends on v 's immediate neighborhood. As a consequence, a potentially interesting avenue for further research is to investigate distributed variants of our approach, possibly with massive parallel architectures in mind [36].

ACKNOWLEDGMENTS

Luca Pepè Sciarria is supported by Spoke 1 “FutureHPC & BigData” of the Italian Research Center on High Performance Computing, Big Data and Quantum Computing (ICSC) funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di “campioni nazionali” di R&S (M4C2-19) - Next Generation EU (NGEU). Luca Becchetti was partially supported by PNRR PE1 Future AI Research (FAIR).

REFERENCES

- [1] Pankaj K Agarwal, Graham Cormode, Zengfeng Huang, Jeff M Phillips, Zhewei Wei, and Ke Yi. 2013. Mergeable summaries. *ACM Transactions on Database Systems (TODS)* 38, 4 (2013), 1–28.
- [2] Charu Aggarwal and Karthik Subbian. 2014. Evolutionary network analysis: A survey. *ACM Computing Surveys (CSUR)* 47, 1 (2014), 1–36.
- [3] Lars Backstrom, Paolo Boldi, Marco Rosa, Johan Ugander, and Sebastiano Vigna. 2012. Four degrees of separation. In *Proceedings of the 4th Annual ACM Web Science Conference* (Evanston, Illinois) (*WebSci '12*). Association for Computing Machinery, New York, NY, USA, 33–42. <https://doi.org/10.1145/2380718.2380723>
- [4] Sachin U. Balvir, Mukesh M. Raghuwanshi, and Purushottam D. Shobhane. 2024. An Overview of Similarity-Based Methods in Predicting Social Network Links: A Comparative Analysis. *IEEE Access* 12 (2024), 120913–120934. <https://doi.org/10.1109/ACCESS.2024.3450506>
- [5] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, D. Sivakumar, and Luca Trevisan. 2002. Counting Distinct Elements in a Data Stream. In *Proceedings of the 6th International Workshop on Randomization and Approximation Techniques (RANDOM '02)*. Springer-Verlag, Berlin, Heidelberg, 1–10.
- [6] Luca Becchetti, Paolo Boldi, Carlos Castillo, and Aristides Gionis. 2008. Efficient semi-streaming algorithms for local triangle counting in massive graphs. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Las Vegas, Nevada, USA) (*KDD '08*). Association for Computing Machinery, New York, NY, USA, 16–24. <https://doi.org/10.1145/1401890.1401898>
- [7] Luca Becchetti, Carlos Castillo, Debora Donato, Ricardo Baeza-Yates, and Stefano Leonardi. 2008. Link analysis for web spam detection. *ACM Transactions on the Web (TWEB)* 2, 1 (2008), 1–42.
- [8] Luca Becchetti, Andrea Clementi, Luciano Gualà, Luca Pepè Sciarria, Alessandro Straziota, and Matteo Stromieri. 2025. Approximate 2-hop neighborhoods on incremental graphs: An efficient lazy approach. *CoRR* abs/2502.19205 (2025). <https://doi.org/10.48550/ARXIV.2502.19205>
- [9] Elisabetta Bergamini, Michele Borassi, Pierluigi Crescenzi, Andrea Marino, and Henning Meyerhenke. 2019. Computing top-k closeness centrality faster in unweighted graphs. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 13, 5 (2019), 1–40.
- [10] Paolo Boldi, Marco Rosa, and Sebastiano Vigna. 2011. HyperANF: approximating the neighbourhood function of very large graphs on a budget. In *Proceedings of the 20th International Conference on World Wide Web* (Hyderabad, India) (*WWW '11*). Association for Computing Machinery, New York, NY, USA, 625–634. <https://doi.org/10.1145/1963405.1963493>
- [11] Paolo Boldi and Sebastiano Vigna. 2014. Axioms for centrality. *Internet Mathematics* 10, 3–4 (2014), 222–262.
- [12] Katy Börner, Jeegar T Maru, and Robert L Goldstone. 2004. The simultaneous evolution of author and paper networks. *Proceedings of the National Academy of Sciences* 101, suppl_1 (2004), 5266–5273.
- [13] A. Broder. 1997. On the Resemblance and Containment of Documents. In *Proceedings of the Compression and Complexity of Sequences 1997 (SEQUENCES '97)*. IEEE Computer Society, USA, 21.
- [14] Andrei Z. Broder. 2000. Identifying and Filtering Near-Duplicate Documents. In *Combinatorial Pattern Matching*, Raffaele Giancarlo and David Sankoff (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–10.
- [15] Andrei Z Broder and Michael Mitzenmacher. 2001. Completeness and robustness properties of min-wise independent permutations. *Random Structures & Algorithms* 18, 1 (2001), 18–30.
- [16] Luciana S. Buriol, Gereon Frahlhing, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Christian Sohler. 2006. Counting triangles in data streams. In *Proceedings of the Twenty-Fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (Chicago, IL, USA) (*PODS '06*). Association for Computing Machinery, New York, NY, USA, 253–262. <https://doi.org/10.1145/1142351.1142388>
- [17] Marc Bury, Chris Schwiigelshohn, and Mara Sorella. 2020. Similarity Search for Dynamic Data Streams. *IEEE Transactions on Knowledge and Data Engineering* 32 (2020), 2241–2253. <https://api.semanticscholar.org/CorpusID:181612911>
- [18] Amit Chakrabarti, Graham Cormode, and Andrew McGregor. 2008. Robust lower bounds for communication and stream computation. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (Victoria, British Columbia, Canada) (*STOC '08*). Association for Computing Machinery, New York, NY, USA, 641–650. <https://doi.org/10.1145/1374376.1374470>
- [19] Qing Chen, Oded Lachish, Sven Helmer, and Michael H. Böhlen. 2022. Dynamic spanning trees for connectivity queries on fully-dynamic undirected graphs. *Proc. VLDB Endow.* 15, 11 (July 2022), 3263–3276. <https://doi.org/10.14778/3551793.3551868>
- [20] Andrea Clementi, Luciano Gualà, Luca Pepè Sciarria, and Alessandro Straziota. 2025. Maintaining k -MinHash Signatures over Fully-Dynamic Data Streams with Recovery. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining, WSDM 2025, Hannover, Germany, March 10-14, 2025*, Wolfgang Nejdl, Sören Auer, Meeyoung Cha, Marie-Francine Moens, and Marc Najork (Eds.). ACM, Hannover, Germany, 79–87. <https://doi.org/10.1145/3701551.3703491>
- [21] Edith Cohen and Haim Kaplan. 2007. Summarizing data using bottom-k sketches. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing* (Portland, Oregon, USA) (*PODC '07*). Association for Computing Machinery, New York, NY, USA, 225–234. <https://doi.org/10.1145/1281100.1281133>
- [22] Reinhard Diestel. 2017. *Graph Theory* (5th ed.). Springer Publishing Company, Incorporated, New York City, United States.
- [23] Devdatt Dubhashi and Alessandro Panconesi. 2009. *Concentration of Measure for the Analysis of Randomized Algorithms* (1st ed.). Cambridge University Press, USA.
- [24] Stanley C. Eisenstat, M. C. Gursky, Martin H. Schultz, and Andrew H. Sherman. 1982. Yale sparse matrix package I: The symmetric codes. *Internat. J. Numer. Methods Engrg.* 18 (1982), 1145–1151. <https://api.semanticscholar.org/CorpusID:120546375>
- [25] Paul Erdős. 1965. On some extremal problems in graph theory. *Israel Journal of Mathematics* 3 (1965), 113–116.
- [26] Shimon Even and Yossi Shiloach. 1981. An On-Line Edge-Deletion Problem. *J. ACM* 28, 1 (1981), 1–4. <https://doi.org/10.1145/322234.322235>
- [27] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. 2005. On graph problems in a semi-streaming model. *Theoretical Computer Science* 348, 2–3 (2005), 207–216.
- [28] Philippe Flajolet and G Nigel Martin. 1985. Probabilistic counting algorithms for data base applications. *Journal of computer and system sciences* 31, 2 (1985), 182–209.
- [29] Linton C Freeman et al. 2002. Centrality in social networks: Conceptual clarification. *Social network: critical concepts in sociology*. Londres: Routledge 1, 3 (2002), 238–263.
- [30] Phillip B. Gibbons and Srikanta Tirathapura. 2001. Estimating simple functions on the union of data streams. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures* (Crete Island, Greece) (*SPAA '01*). Association for Computing Machinery, New York, NY, USA, 281–291. <https://doi.org/10.1145/378580.378687>
- [31] Kathrin Hanauer, Monika Henzinger, and Christian Schulz. 2022. Recent Advances in Fully Dynamic Graph Algorithms – A Quick Reference Guide. *ACM J. Exp. Algorithmics* 27 (2022), 1.11:1–1.11:45. <https://doi.org/10.1145/3555806>
- [32] Kathrin Hanauer, Monika Henzinger, and Christian Schulz. 2022. Recent Advances in Fully Dynamic Graph Algorithms – A Quick Reference Guide. *ACM J. Exp. Algorithmics* 27, Article 1.11 (Dec. 2022), 45 pages. <https://doi.org/10.1145/3555806>
- [33] Peng Jia, Pinghui Wang, Jing Tao, and Xiaohong Guan. 2019. A Fast Sketch Method for Mining User Similarities Over Fully Dynamic Graph Streams. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE Computer Society, USA, 1682–1685. <https://doi.org/10.1109/ICDE.2019.00172>
- [34] Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. 2014. Approximating matching size from random streams. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms* (Portland, Oregon) (*SODA '14*). Society for Industrial and Applied Mathematics, USA, 734–751.
- [35] Matthieu Latapy, Clémence Magnien, and Nathalie Del Vecchio. 2008. Basic notions for the analysis of large two-mode networks. *Social Networks* 30, 1 (2008), 31–48. <https://doi.org/10.1016/j.socnet.2007.04.006>
- [36] Kyong-Ha Lee, Yoon-Joon Lee, Hyunsik Choi, Yon Dohn Chung, and Bongki Moon. 2012. Parallel data processing with MapReduce: a survey. *AcM SIGMOD record* 40, 4 (2012), 11–20.
- [37] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Paris, France) (*KDD '09*). Association for Computing Machinery, New York, NY, USA, 497–506. <https://doi.org/10.1145/1557019.1557077>
- [38] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densification and shrinking diameters. *ACM transactions on Knowledge Discovery from Data (TKDD)* 1, 1 (2007), 2–es.
- [39] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [40] David Liben-Nowell and Jon Kleinberg. 2003. The link prediction problem for social networks. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management* (New Orleans, LA, USA) (*CIKM '03*). Association for Computing Machinery, New York, NY, USA, 556–559. <https://doi.org/10.1145/956863.956972>
- [41] Andrew McGregor. 2014. Graph stream algorithms: a survey. *ACM SIGMOD Record* 43, 1 (2014), 9–20.
- [42] Morteza Monemizadeh, S. Muthukrishnan, Pan Peng, and Christian Sohler. 2017. Testable Bounded Degree Graph Properties Are Random Order Streamable. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017) (Leibniz International Proceedings in Informatics (LIPIcs))*, Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl (Eds.), Vol. 80. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 131:1–131:14. <https://doi.org/10.4230/LIPIcs.ICALP.2017.131>

- [43] Georgios A Pavlopoulos, Panagiota I Kontou, Athanasia Pavlopoulou, Costas Bouyioukos, Evripides Markou, and Pantelis G Bagos. 2018. Bipartite graphs in systems biology and medicine: a survey of methods and applications. *GigaScience* 7, 4 (2018), giy014.
- [44] Pan Peng and Christian Sohler. 2018. Estimating graph parameters from random order streams. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, Louisiana) (SODA '18). Society for Industrial and Applied Mathematics, USA, 2449–2466.
- [45] Mihai Pătraşcu and Mikkel Thorup. 2012. The Power of Simple Tabulation Hashing. *J. ACM* 59, 3, Article 14 (June 2012), 50 pages. <https://doi.org/10.1145/2220357.2220361>
- [46] Yannick Rochat. 2009. Closeness Centrality Extended to Unconnected Graphs: the Harmonic Centrality Index. In *ASNA, Zürich, August 26-28, 2009*. Institute of Applied Mathematics, Switzerland, 1–15.
- [47] Liam Roditty and Uri Zwick. 2011. On Dynamic Shortest Paths Problems. *Algorithmica* 61, 2 (2011), 389–401. <https://doi.org/10.1007/S00453-010-9401-5>
- [48] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The network data repository with interactive graph analytics and visualization. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, Austin, Texas, 4292–4293.
- [49] Ahmet Erdem Sarıyüce and Ali Pinar. 2018. Peeling Bipartite Networks for Dense Subgraph Discovery. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) (WSDM '18). Association for Computing Machinery, New York, NY, USA, 504–512. <https://doi.org/10.1145/3159652.3159678>
- [50] Robert Endre Tarjan. 1985. Amortized Computational Complexity. *SIAM Journal on Algebraic Discrete Methods* 6, 2 (1985), 306–318. <https://doi.org/10.1137/0606031> arXiv:<https://doi.org/10.1137/0606031>
- [51] Bimal Viswanath, Alan Mislove, Meeyoung Cha, and Krishna P. Gummadi. 2009. On the evolution of user interaction in Facebook. In *Proceedings of the 2nd ACM Workshop on Online Social Networks* (Barcelona, Spain) (WOSN '09). Association for Computing Machinery, New York, NY, USA, 37–42. <https://doi.org/10.1145/1592665.1592675>
- [52] Xu-Wen Wang, Lorenzo Madeddu, Kerstin Spirohn, Leonardo Martini, Adriano Fazzzone, Luca Becchetti, Thomas P Wytock, István A Kovács, Olivér M Balogh, Bettina Benczik, et al. 2023. Assessment of community efforts to advance network-based prediction of protein–protein interactions. *Nature communications* 14, 1 (2023), 1582.
- [53] Qingjun Xiao, Shiwei Yang, Panpan Li, Kangying Li, and Lin Wen. 2024. Multi-Resolution Odd Sketch for Mining Extended Jaccard Similarity of Dynamic Streaming sets. *IEEE Transactions on Network Science and Engineering* 11, 3 (2024), 2399–2414. <https://doi.org/10.1109/TNSE.2023.3275809>
- [54] Jaewon Yang and Jure Leskovec. 2012. Defining and evaluating network communities based on ground-truth. In *Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics* (Beijing, China) (MDS '12). Association for Computing Machinery, New York, NY, USA, Article 3, 8 pages. <https://doi.org/10.1145/2350190.2350193>
- [55] Ahmad Zareie and Rizos Sakellariou. 2020. Similarity-based link prediction in social networks using latent relationships between the users. *Scientific reports* 10, 1 (2020), 20137.
- [56] Fangyuan Zhang and Sibor Wang. 2022. Effective indexing for dynamic structural graph clustering. *Proceedings of the VLDB Endowment* 15, 11 (2022), 2908–2920.
- [57] Tao Zhou, Yan-Li Lee, and Guannan Wang. 2021. Experimental analyses on 2-hop-based and 3-hop-based link prediction algorithms. *Physica A: Statistical Mechanics and its Applications* 564 (2021), 125532.