



HAL
open science

A Streaming Algorithm for Graph Clustering

Alexandre Hollocou, Julien Maudet, Thomas Bonald, Marc Lelarge

► **To cite this version:**

Alexandre Hollocou, Julien Maudet, Thomas Bonald, Marc Lelarge. A Streaming Algorithm for Graph Clustering. NIPS 2017 - Workshop on Advances in Modeling and Learning Interactions from Complex Data, Dec 2017, Long Beach, United States. pp.1-10. hal-01639506v1

HAL Id: hal-01639506

<https://hal.science/hal-01639506v1>

Submitted on 20 Nov 2017 (v1), last revised 8 Dec 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Streaming Algorithm for Graph Clustering

Alexandre Hollocou

INRIA
Paris, France
alexandre.hollocou@inria.fr

Julien Maudet

Ecole Polytechnique
Palaiseau, France
julien.maudet@polytechnique.edu

Thomas Bonald

Telecom-Paristech
Paris, France
thomas.bonald@telecom-paristech.fr

Marc Lelarge

INRIA-ENS
Paris, France
marc.lelarge@ens.fr

Abstract

We introduce a novel algorithm to perform graph clustering in the edge streaming setting. In this model, the graph is presented as a sequence of edges that can be processed strictly once. Our streaming algorithm has an extremely low memory footprint as it stores only three integers per node and does not keep any edge in memory. We provide a theoretical justification of the design of the algorithm based on the *modularity* function, which is a usual metric to evaluate the quality of a graph partition. We perform experiments on massive real-life graphs ranging from one million to more than one billion edges and we show that this new algorithm runs more than ten times faster than existing algorithms and leads to similar or better detection scores on the largest graphs.

1 Introduction

1.1 Motivations

Graphs arise in a wide range of fields from biology [23] to social media [20] or web analysis [10][25]. In most of these graphs, we observe groups of nodes that are densely connected between each other and sparsely connected to the rest of the graph. One of the most fundamental problems in the study of such graphs consists in identifying these dense clusters of nodes. This problem is commonly referred to as *graph clustering* or *community detection*.

A major challenge for community detection algorithms is their ability to process very large graphs that are commonly observed in numerous fields. For instance, social networks have typically millions of nodes and billions of edges (e.g. Friendster [20]). Many algorithms have been proposed during the last ten years, using various techniques ranging from combinatorial optimization to spectral analysis [14]. Most of them fail to scale to such large real-life graphs [27] and require the whole graph to be stored in memory, which often represents a heavy constraint in practice. Streaming the edges is a natural way to handle such massive graphs. In this setting, the entire graph is not stored but processed edge by edge [19]. Note that the streaming approach is particularly relevant in most real-life applications where graphs are fundamentally dynamic and edges naturally arrive in a streaming fashion.

1.2 Contributions

In this paper, we introduce a novel approach based on *edge streams* to detect communities in graphs. The algorithm processes each edge strictly once. When the graph is a multi-graph, in the sense that

two nodes may be connected by more than one edge, these edges are streamed independently. The algorithm only stores three integers for each node: its current community index, its current degree (i.e. the number of adjacent edges that have already been processed), and the current community volume (i.e. the sum of the degrees of all nodes in the community). Hence, the time complexity of the algorithm is linear in the number of edges and its space complexity is linear in the number of nodes. In the experimental evaluation of the algorithm, we show that this streaming algorithm is able to handle massive graphs [35] with low execution time and memory consumption.

The algorithm takes only one integer parameter v_{\max} and, for each arriving edge (i, j) of the stream, it uses a simple decision strategy based on this parameter and the volumes of the communities of nodes i and j . We provide a theoretical analysis that justifies the form of this decision strategy using the so-called *modularity* of the clustering. Modularity, that has been introduced by the physics community [22], is one of the most widely used quality function for graph clustering. It measures the quality of a given partition based on the comparison between the number of edges that are observed in each cluster and the number of edges that would be observed if the edges were randomly distributed. In our analysis, we show that, under certain assumptions, the processing of each new edge by our algorithm leads to an increase in modularity.

1.3 Related work

A number of algorithms have been developed for detecting communities in graphs [11]. Many rely on the optimization of some objective function that measures the quality of the detected communities. Modularity and other metrics, like conductance, out-degree fraction and the clustering coefficient [35] have been used with success. Other popular methods include spectral clustering [29][31], clique percolation [23], statistical inference [16], random walks [26][32] and matrix factorization [34]. These techniques have proved to be efficient but are often time-consuming and fail to scale to large graphs [27].

The streaming approach has drawn considerable interest in network analysis over the last decade. Within the data stream model, massive graphs with potentially billions of edges can be processed without being stored in memory [19]. A lot of algorithm have been proposed for different problems that arise in large graphs, such as counting subgraphs [3][6], computing matchings [13][9], finding the minimum spanning tree [7][30] or graph sparsification [4]. Different types of data streams can be considered: *insert-only streams*, where the stream is the unordered sequence of the network edges, or *dynamic graph streams*, where edges can both be added or deleted. Many streaming algorithms rely on *graph sketches* which store the input in a memory-efficient way and are updated at each step [1].

In this paper, we use the streaming setting to define a novel community detection algorithm. We use an insert-only edge streams and define a minimal sketch, by storing only three integers per node.

1.4 Paper outline

The rest of the paper is organized as follows. We first describe our streaming algorithm in Section 2. A theoretical analysis of this algorithm is presented in Section 3. In Section 4, we evaluate experimentally the performance of our approach on real-life graphs and compare it to state-of-the-art algorithms. Section 5 concludes the paper.

2 A streaming algorithm for community detection

In this section, we define a novel streaming algorithm for community detection in graphs.

2.1 Streaming setting

We are given an undirected and unweighted multi-graph $G(V, E)$ where V is the set of vertices and E is a multi-set of edges (i.e. an edge (i, j) can appear multiple times in E). We use n to denote the number of nodes and m the number of edges. We use w_{ij} to denote the number of edges between i and j (and we set $w_{ij} = 0$ if $(i, j) \notin E$). We assume that there is no self-loop, that is $w_{ii} = 0$ for all $i \in V$. We use w_i to denote the degree $\sum_j w_{ij}$ of node i , and $w = \sum_i w_i = 2m$ to denote the weight of the graph, corresponding to the total degree. Given a set of nodes C , we use $\text{Vol}(C) = \sum_{i \in C} w_i$ to denote the volume of C .

We consider the following streaming framework: we are given a stream $S = (e_1, \dots, e_m)$, which is an order sequence of the multi-set E . Note that each edge $e = (i, j) \in E$ appears exactly w_{ij} times in S .

2.2 Intuition

Although there is no universal definition of what a community is, most existing algorithms rely on the principle that nodes tend to be more connected within a community than across communities. Hence, if we pick uniformly at random an edge e in E , this edge is more likely to link nodes of the same community (i.e., e is an *intra-community* edge), than nodes from distinct communities (i.e., e is an *inter-community* edge). Equivalently, if we assume that edges arrive in a random order, we expect many intra-community edges to arrive before the inter-community edges.

This observation is used to design a streaming algorithm. For each arriving edge (i, j) , the algorithm places i and j in the same community if the edge arrives *early* (intra-community edge) and splits the nodes in distinct communities otherwise (inter-community edge). In this formulation, the notion of an *early* edge is of course critical. In the proposed algorithm, we consider that an edge (i, j) arrives *early* if the current volumes of the communities of nodes i and j , accounting for previously arrived edges only, is low.

More formally, the algorithm considers successively each edge of the stream $S = (e_1, e_2, \dots, e_m)$. Each node is initially in its own community. At time t , a new edge $e_t = (i, j)$ arrives and the algorithm performs one of the following actions: (a) i joins the community of j ; (b) j joins the community of i ; (c) no action.

The choice of the action depends on the *updated* community volumes $\text{Vol}(C(i))$ and $\text{Vol}(C(j))$ of the communities of i and j , $C(i)$ and $C(j)$, i.e., the volumes computed using the edges e_1, \dots, e_t . If $\text{Vol}(C(i))$ or $\text{Vol}(C(j))$ is greater than a given threshold v_{\max} , then we do nothing; otherwise, the node belonging to the smallest community (in volume) joins the community of the other node and the volumes are updated.

2.3 Algorithm

We define our streaming algorithm in Algorithm 1. It takes the list of edges of the graph and one integer parameter $v_{\max} \geq 1$. The algorithm uses three dictionaries d , c and v , initialized with default value 0. At the end of the algorithm, d_i is the degree of node i , c_i the community of node i , and v_k is the volume of community k . When an edge with an unknown node arrives, let say i , we give this node a new community index, $c_i \leftarrow k$, and increment the index variable k (which is initialized with 1). For each new edge $e = (i, j)$, the degrees of i and j and the volumes of communities c_i and c_j are updated. Then, if these volumes are both lower than the threshold parameter v_{\max} , the node in the community with the lowest volume joins the community of the other node. Otherwise, the communities remain unchanged.

Observe that, in case of equality $v_{c_i} = v_{c_j} \leq v_{\max}$, j joins the community of i . Of course, this choice is arbitrary and can be made random (e.g., i joins the community of j with probability $1/2$ and j joins the community of i with probability $1/2$).

2.4 Complexity

The main loop is linear in the number of edges in the stream. Thus, the time complexity of the algorithm is linear in m .

Concerning the space complexity, we only use three dictionaries of integers d , c and v , of size n . Hence, the space complexity of the algorithm is $3n \cdot \text{sizeof}(\text{int}) = O(n)$. Note that the algorithm does not need to store the list of edges in memory, which is the main benefit of the streaming approach. To implement dictionaries with default value 0 in practice, we can use classic dictionaries or maps, and, when an unknown key is requested, set the value relative to this key to 0. Note that, in Python, the `defaultdict` structure already allows us to exactly implement dictionaries with 0 as a default value.

Algorithm 1 Streaming algorithm for clustering graph nodes

Require: Stream of edges S and parameter $v_{\max} \geq 1$

- 1: $d, v, c \leftarrow$ dictionaries initialized with default value 0
- 2: $k \leftarrow 1$ (*new community index*)
- 3: **for** $(i, j) \in S$ **do**
- 4: **if** $c_i = 0$ **then** $c_i \leftarrow k$ and $k \leftarrow k + 1$
- 5: **end if**
- 6: **if** $c_j = 0$ **then** $c_j \leftarrow k$ and $k \leftarrow k + 1$
- 7: **end if**
- 8: $d_i \leftarrow d_i + 1$ and $d_j \leftarrow d_j + 1$ (*update degrees*)
- 9: $v_{c_i} \leftarrow v_{c_i} + 1$ and $v_{c_j} \leftarrow v_{c_j} + 1$ (*update community volumes*)
- 10: **if** $v_{c_i} \leq v_{\max}$ and $v_{c_j} \leq v_{\max}$ **then**
- 11: **if** $v_{c_i} \leq v_{c_j}$ **then** (*i joins the community of j*)
- 12: $v_{c_j} \leftarrow v_{c_j} + d_i$
- 13: $v_{c_i} \leftarrow v_{c_i} - d_i$
- 14: $c_i \leftarrow c_j$
- 15: **else** (*j joins the community of i*)
- 16: $v_{c_i} \leftarrow v_{c_i} + d_j$
- 17: $v_{c_j} \leftarrow v_{c_j} - d_j$
- 18: $c_j \leftarrow c_i$
- 19: **end if**
- 20: **end if**
- 21: **end for**
- 22: **return** $(c_i)_{i \in V}$

2.5 Parameter setting

Note that the algorithm can be run once with multiple values of parameter v_{\max} . In this case, only arrays c and v need to be duplicated for each value of v_{\max} . In this multi-parameter setting, we obtain multiple results $(c^a)_{1 \leq a \leq A}$ at the end of the algorithm, where A is the number of distinct values for the parameter v_{\max} . Then, the best c^a can be selected by computing quality metrics that only use dictionaries c^a and v^a . In particular, we do not want to use metrics that requires the knowledge of the input graph. For instance, common metrics [35] like entropy $H(v) = -\sum_k \frac{v_k}{w} \log\left(\frac{v_k}{w}\right)$ or average density $D(c, v) = \sum_{k: C_k \neq \emptyset} \frac{1}{|P|} \frac{v_k}{|C_k|(|C_k|-1)}$, where C_k is the set of nodes in community k and P is the set of all non-empty communities, can be easily computed from each pair of dictionaries (v^a, c^a) , and be used to select the best result c^a for $a = 1, \dots, A$. Note that modularity cannot be used here as its computation requires the knowledge of the whole graph.

3 Theoretical analysis

In this section, we analyze the modularity optimization problem in the edge-streaming setting and qualitatively justify the conditions on community volumes used in Algorithm 1. Complete proofs are available as supplementary material on GitHub¹.

3.1 Modularity optimization in the streaming setting

Modularity is a quality metric that is widely used in graph clustering [22]. Given a partition P of the nodes, modularity is defined as

$$Q = \frac{1}{w} \sum_{i \in V} \sum_{j \in V} \left(w_{ij} - \frac{w_i w_j}{w} \right) \delta(i, j)$$

where $\delta(i, j) = 1$ if i and j belongs to the same community $C \in P$ and 0 otherwise. Modularity can be seen as the difference between two probabilities,

$$Q = \mathbb{P}[(i, j) \sim E : C(i) = C(j)] - \mathbb{P}[(i, j) \sim \mathcal{N} : C(i) = C(j)],$$

¹<https://github.com/ahollocou/graph-streaming>

where $\mathbb{P}[(i, j) \sim E : C(i) = C(j)]$ corresponds to the probability to choose an edge of G uniformly at random between two nodes of the same community, and $\mathbb{P}[(i, j) \sim \mathcal{N} : C(i) = C(j)]$ to the probability to choose an edge between two nodes from the same community in the so-called *null model* \mathcal{N} , where an edge (i, j) is chosen with a probability proportional to the degrees of i and j . A classic approach to cluster graphs consists in finding a partition that maximizes Q . Many algorithms have been proposed to perform this task [5, 21] but, to the best of our knowledge, none can be applied to our streaming setting.

The modularity can be rewritten as

$$Q = \frac{1}{w} \left[\sum_{i \in V} \sum_{j \in V} w_{ij} \delta(i, j) - \sum_{C \in \mathcal{P}} \frac{\text{Vol}(C)^2}{w} \right].$$

In our streaming setting, we are given a stream $S = (e_1, \dots, e_m)$ of edges such that edge (i, j) appears w_{ij} times in S . We consider the situation where t edges e_1, \dots, e_t from the stream $S = (e_1, \dots, e_m)$ have already arrived, and where we have computed a partition $P_t = (C_1, \dots, C_K)$ of the graph. We define S_t as $S_t = \{e_1, \dots, e_t\}$, and Q_t as

$$Q_t = \sum_{C \in P_t} \left[2\text{Int}_t(C) - \frac{(\text{Vol}_t(C))^2}{w} \right]$$

where $\text{Int}_t(C) = \sum_{(i, j) \in S_t} 1_{i \in C} 1_{j \in C}$ and $\text{Vol}_t(C) = \sum_{(i, j) \in S_t} (1_{i \in C} + 1_{j \in C})$. Note that there is no normalization factor $1/w$ in the definition of Q_t as it has no impact on the optimization problem.

We do not store the edges of S_t but we assume that we have kept updated node degrees $w_t(i) = \sum_{(i', j') \in S_t} (1_{i'=i} + 1_{j'=i})$ and community volumes $\text{Vol}_t(C_k)$ in a streaming fashion as shown in Algorithm 1. We consider the situation where a new edge $e_{t+1} = (i, j)$ arrives. We want to make a decision that maximizes Q_{t+1} .

3.2 Streaming decision

We can express Q_{t+1} in function of Q_t as stated in Lemma 1.

Lemma 1. *If $e_{t+1} = (i, j)$ and if $P_{t+1} = P_t$, Q_{t+1} can be expressed in function of Q_t as follows*

$$Q_{t+1} = Q_t + 2 \left[\delta(i, j) - \frac{\text{Vol}_t(C(i)) + \text{Vol}_t(C(j)) + 1 + \delta(i, j)}{w} \right]$$

where $C(v)$ denotes the community of v in P_t , and $\delta(i, j) = 1$ if i and j belongs to the same community of P_t and 0 otherwise.

We want to update the community membership of i or j with one of the following actions: (a) i joins the community of j ; (b) j joins the community of i ; (c) i and j stays in their respective communities. We consider the case where nodes i and j belongs to distinct communities of P_t , since all three actions are identical if i and j belong to the same community. We want to choose the action that maximizes Q_{t+1} , but we have a typical streaming problem where we cannot evaluate the impact of action (a) or (b) on Q_t but only on the term that comes from the new edge e_{t+1} .

Let us consider action (a), where i joins the community of j . We can assume that $\text{Vol}_t(C(i)) \leq \text{Vol}_t(C(j))$ without loss of generality (otherwise we can swap i and j). We are interested in $\Delta Q_{t+1} = Q_{t+1}^{(a)} - Q_{t+1}^{(c)}$, which is the variation of Q_{t+1} between the state where i and j are in their own communities and the state where i has joined $C(j)$. We have $\Delta Q_{t+1} = \Delta Q_t + 2[1 - (\text{Vol}_t(C(j)) - \text{Vol}_t(C(i)) + 2w_i(t) + 1)/w]$ where ΔQ_t is the variation of Q_t . Lemma 2 gives us an expression for this variation.

Lemma 2.

$$\Delta Q_t = Q_t^{(a)} - Q_t^{(c)} = 2 \left[L_t(i, C(j)) - L_t(i, C(i)) - \frac{(w_t(i))^2}{w} \right]$$

where

$$L_t(i, C) = \sum_{(i', j') \in S_t} \left[1_{i' \in C} \left(1_{j'=i} - \frac{w_t(i)}{w} \right) + 1_{j' \in C} \left(1_{i'=i} - \frac{w_t(i)}{w} \right) \right].$$

We define $l_t(i, C)$ as $l_t(i, C) = L_t(i, C)/\text{Vol}_t(C)$. Then, we can easily show that $l_t(i, C) \in [-1, 1]$ and $\mathbb{E}[l_t(i, C)] = 0$ if edges of S_t follow the null model presented above; $l_t(i, C)$ can be seen as a *normalized degree of attachment of node i to community C* in the edge stream S_t .

Lemma 1 and 2 gives us a sufficient condition presented in Theorem 1 in order to have a positive variation ΔQ_{t+1} of the modularity when i joins $C(j)$.

Theorem 1. *If $\text{Vol}_t(C(i)) \leq \text{Vol}_t(C(j))$, then:*

$$\text{Vol}_t(C(j)) \leq v_t(i, j) \implies \Delta Q_{t+1} \geq 0$$

where

$$v_t(i, j) = \begin{cases} \frac{1 - (w_t(i) + 1)^2 / w}{l_t(i, C(i)) - l_t(i, C(j))} & \text{if } l_t(i, C(i)) \neq l_t(i, C(j)) \\ +\infty & \text{otherwise.} \end{cases}$$

Thus, we see that, if we have $v_{\max} \leq v_t(i, j)$, then the strategy used by Algorithm 1 leads to an increase in modularity. In the general case, we cannot control terms $l_t(i, C(i))$ and $l_t(i, C(j))$, but, in most cases, we expect the degree of attachment of i to $C(i)$ to be upper-bounded by some constant $\tau_1 < 1$, the degree of attachment of i to $C(j)$ to be higher than some $\tau_2 > 0$, and the term $(w_t(i) + 1)^2 / w$ to be smaller than a constant $\epsilon \ll 1$. Then, the condition on v_{\max} becomes:

$$v_{\max} \leq \frac{1 - \epsilon}{\tau_1 - \tau_2}.$$

This justifies the design of the algorithm, with the decision of joining one community or the other based on the community volumes.

4 Experimental results

4.1 Datasets

We use real-life graphs provided by the Stanford Social Network Analysis Project (SNAP [35]) for the experimental evaluation of our new algorithm. These datasets include ground-truth community memberships that we use to measure the quality of the detection. We consider datasets of different natures. **Social networks**: The YouTube, LiveJournal, Orkut and Friendster datasets correspond to social networks [2][20] where nodes represent users and edges connect users who have a friendship relation. In all these networks, users can create groups that are used as ground-truth communities in the dataset definitions. **Co-purchasing network**: The Amazon dataset corresponds to a product co-purchasing network [18]. The nodes of the graph represent Amazon products and the edges correspond to frequently co-purchased products. The ground-truth communities are defined as the product categories. **Co-citation network**: The DBLP dataset corresponds to a scientific collaboration network [2]. The nodes of the graph represent the authors and the edges the co-authorship relations. The scientific conferences are used as ground-truth communities.

The size of these graphs ranges from approximately one million edges to more than one billion edges. It enables us to test the ability of our algorithm to scale to very large graphs. The characteristics of these datasets can be found in Table 1.

4.2 Benchmark algorithms

For assessing the performance of our streaming algorithm we use a wide range of state-of-the-art but non-streaming algorithms that are based on various approaches. **SCD (S)** partitions the graph by maximizing the WCC, which is a community quality metric based on triangle counting [27]. **Louvain (L)** is based on the optimization of the well-known modularity metric [5]. **Infomap (I)** splits the network into modules by compressing the information flow generated by random walks [28]. **Walktrap (W)** uses random walks to estimate the similarity between nodes, which is then used to cluster the network [26]. **OSLOM (O)** partitions the network by locally optimizing a fitness function which measures the statistical significance of a community [17]. In the data tables, we use **STR** to refer to our streaming algorithm.

4.3 Performance metrics and benchmark setup

We use two metrics for the performance evaluation of the selected algorithms. The first is the *average F1-score* [34][27] which corresponds to the harmonic mean of precision and recall. The second metric is the *Normalized Mutual Information* (NMI), which is based on the mutual entropy between indicator functions for the communities [15].

The experiments were performed on EC2 instances provided by Amazon Web Services of type m4.4xlarge with 64 GB of RAM, 100 GB of disk space, 16 virtual CPU with Intel Xeon Broadwell or Haswell and Ubuntu Linux 14.04 LTS.

Our algorithm is implemented in C++ and the source code can be found on GitHub². For the other algorithms, we used the C++ implementations provided by the authors, that can be found on their respective websites. Finally, all the scoring functions were implemented in C++. We used the implementation provided by the authors of [15] for the NMI and the implementation provided by the authors of SCD [27] for the F1-Score.

4.4 Benchmark results

Execution time

We compare the execution times of the different algorithms on SNAP graphs in Table 1. The entries that are not reported in the table corresponds to algorithms that returned execution errors or algorithms with execution times exceeding 6 hours. In our experiments, only SCD, except from our algorithm, was able to run on all datasets. The fastest algorithms in our benchmarks are SCD and Louvain and we observe that they run more than ten times slower than our streaming algorithm. More precisely, our streaming algorithm runs in less than 50ms on the Amazon and DBLP graphs, which contain millions of edges, and less than 5 minutes on the largest network, Friendster, that has more than one billion edges. In comparison, it takes seconds for SCD and Louvain to detect communities on the smallest graphs, and several hours to run on Friendster. Table 1 shows the execution times of all the algorithms with respect to the number of edges in the network. We remark that there is more than one order of magnitude between our algorithm and the other algorithms.

In order to compare the execution time of our algorithm with a minimal algorithm that only reads the list of edges without doing any additional operation, we measured the run time of the Unix command `cat` on the largest dataset, Friendster. `cat` reads the edge file sequentially and writes each line corresponding to an edge to standard output. In our experiments, the command `cat` takes 152 seconds to read the list of edges of the Friendster dataset, whereas our algorithm processes this network in 241 seconds. That is to say, reading the edge stream is only twice faster than the execution of our streaming algorithm.

	$ V $	$ E $	S	L	I	W	O	STR
Amazon	334,863	925,872	1.84	2.85	31.8	261	1038	0.05
DBLP	317,080	1,049,866	1.48	5.52	27.6	1785	1717	0.05
YouTube	1,134,890	2,987,624	9.96	11.5	150	-	-	0.14
LiveJournal	3,997,962	34,681,189	85.7	206	-	-	-	2.50
Orkut	3,072,441	117,185,083	466	348	-	-	-	8.67
Friendster	65,608,366	1,806,067,135	13464	-	-	-	-	241

Table 1: SNAP dataset sizes and execution times in seconds

Memory consumption

We measured the memory consumption of streaming algorithm and compared it to the memory that is needed to store the list of the edges for each network, which is a lower bound of the memory consumption of the other algorithms. We use 64-bit integers to store the node indices. The memory needed to represent the list of edges is 14,8 MB for the smallest network, Amazon, and 28,9 GB for the largest one, Friendster. In comparison, our algorithm consumes 8,1 MB on Amazon and only 1,6 GB on Friendster.

²<https://github.com/ahollocou/graph-streaming>

	F1-Score						NMI					
	S	L	I	W	O	STR	S	L	I	W	O	STR
Ama.	0.39	0.47	0.30	0.39	0.47	0.38	0.16	0.24	0.16	0.26	0.23	0.12
DBLP	0.30	0.32	0.10	0.22	0.35	0.28	0.15	0.14	0.01	0.10	0.15	0.10
YT	0.23	0.11	0.02	-	-	0.26	0.10	0.04	0.00	-	-	0.13
LiveJ.	0.19	0.08	-	-	-	0.28	0.05	0.02	-	-	-	0.09
Orkut	0.22	0.19	-	-	-	0.44	0.22	0.19	-	-	-	0.24
Friend.	0.10	-	-	-	-	0.19	-	-	-	-	-	-

Table 2: Average F1 Scores and NMI

Detection scores

Table 2 shows the Average F1-score and NMI of the algorithms on the SNAP datasets. Note that the NMI on the Friendster dataset is not reported in the table because the scoring program used for its computation [15] cannot handle the size of the output on this dataset. While Louvain and OSLOM clearly outperform our algorithm on Amazon and DBLP (at the expense of longer execution times), our streaming algorithm shows similar performance as SCD on YouTube and much better performance than SCD and Louvain on LiveJournal, Orkut and Friendster (the other algorithms do not run these datasets). Thus our algorithm does not only run much faster than the existing algorithms but the quality of the detected communities is also better than that of the state-of-the-art algorithms for very large graphs.

5 Conclusion and future work

We introduced a new algorithm for the problem of graph clustering in the edge streaming setting. In this setting, the input data is presented to the algorithm as a sequence of edges that can be examined only once. Our algorithm only stores three integers per node and requires only one integer parameter v_{\max} . It runs more than 10 times faster than state-of-the-art algorithms such as Louvain and SCD and shows better detection scores on the largest graphs. Such an algorithm is extremely useful in many applications where massive graphs arise. For instance, the web graph contains around 10^{10} nodes which is much more than in the Friendster dataset.

We analyzed the adaptation of the popular modularity problem to the streaming setting. Theorem 1 justifies the nature of the condition on the volumes of the communities of nodes i and j for each new edge (i, j) , which is the core of Algorithm 1.

It would be interesting for future work to perform further experiments. In particular the ability of the algorithm to handle evolving graphs could be evaluated on dynamic datasets [24] and compared to existing approaches [12][8]. Note that, in the dynamic network settings, modifications to the algorithm design could be made to handle events such as edge deletions.

Finally, our algorithm only returns disjoint communities, whereas, in many real graphs, overlaps between communities can be observed [15]. An important research direction would consist in adapting approach to overlapping community detection and compare it to existing approaches [33][34].

References

- [1] K. J. Ahn, S. Guha, and A. McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 5–14. ACM, 2012.
- [2] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan. Group formation in large social networks: membership, growth, and evolution. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 44–54. ACM, 2006.
- [3] Z. Bar-Yossef, R. Kumar, and D. Sivakumar. Reductions in streaming algorithms, with an application to counting triangles in graphs. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 623–632. Society for Industrial and Applied Mathematics, 2002.

- [4] A. A. Benczúr and D. R. Karger. Approximating st minimum cuts in $\tilde{O}(n^2)$ time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 47–55. ACM, 1996.
- [5] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008.
- [6] L. S. Buriol, G. Frahling, S. Leonardi, A. Marchetti-Spaccamela, and C. Sohler. Counting triangles in data streams. In *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 253–262. ACM, 2006.
- [7] M. Elkin and J. Zhang. Efficient algorithms for constructing $(1+\epsilon, \beta)$ -spanners in the distributed and streaming models. *Distributed Computing*, 18(5):375–385, 2006.
- [8] A. Epasto, S. Lattanzi, and M. Sozio. Efficient densest subgraph computation in evolving graphs. In *Proceedings of the 24th International Conference on World Wide Web*, pages 300–310. ACM, 2015.
- [9] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. On graph problems in a semi-streaming model. *Theoretical Computer Science*, 348(2-3):207–216, 2005.
- [10] G. W. Flake, S. Lawrence, and C. L. Giles. Efficient identification of web communities. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 150–160. ACM, 2000.
- [11] S. Fortunato. Community detection in graphs. *Physics reports*, 486(3):75–174, 2010.
- [12] L. Gauvin, A. Panisson, and C. Cattuto. Detecting the community structure and activity patterns of temporal networks: a non-negative tensor factorization approach. *PloS one*, 9(1):e86028, 2014.
- [13] A. Goel, M. Kapralov, and S. Khanna. On the communication and streaming complexity of maximum bipartite matching. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 468–485. SIAM, 2012.
- [14] A. Lancichinetti and S. Fortunato. Community detection algorithms: a comparative analysis. *Physical review E*, 80(5):056117, 2009.
- [15] A. Lancichinetti, S. Fortunato, and J. Kertész. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics*, 11(3):033015, 2009.
- [16] A. Lancichinetti, F. Radicchi, and J. J. Ramasco. Statistical significance of communities in networks. *Physical Review E*, 81(4):046110, 2010.
- [17] A. Lancichinetti, F. Radicchi, J. J. Ramasco, and S. Fortunato. Finding statistically significant communities in networks. *PloS one*, 6(4):e18961, 2011.
- [18] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.
- [19] A. McGregor. Graph stream algorithms: a survey. *ACM SIGMOD Record*, 43(1):9–20, 2014.
- [20] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42. ACM, 2007.
- [21] M. E. Newman. Fast algorithm for detecting community structure in networks. *Physical review E*, 69(6):066133, 2004.
- [22] M. E. Newman. Modularity and community structure in networks. *Proceedings of the national academy of sciences*, 103(23):8577–8582, 2006.
- [23] G. Palla, I. Derényi, I. Farkas, and T. Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.

- [24] P. Panzarasa, T. Opsahl, and K. M. Carley. Patterns and dynamics of users' behavior and interaction: Network analysis of an online community. *Journal of the American Society for Information Science and Technology*, 60(5):911–932, 2009.
- [25] R. Pastor-Satorras and A. Vespignani. *Evolution and structure of the Internet: A statistical physics approach*. Cambridge University Press, 2007.
- [26] P. Pons and M. Latapy. Computing communities in large networks using random walks. In *International Symposium on Computer and Information Sciences*, pages 284–293. Springer, 2005.
- [27] A. Prat-Pérez, D. Dominguez-Sal, and J.-L. Larriba-Pey. High quality, scalable and parallel community detection for large real graphs. In *Proceedings of the 23rd international conference on World wide web*, pages 225–236. ACM, 2014.
- [28] M. Rosvall and C. T. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123, 2008.
- [29] D. A. Spielman and S.-H. Teng. Spectral partitioning works: Planar graphs and finite element meshes. *Linear Algebra and its Applications*, 421(2-3):284–305, 2007.
- [30] R. E. Tarjan. *Data structures and network algorithms*. SIAM, 1983.
- [31] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [32] J. J. Whang, D. F. Gleich, and I. S. Dhillon. Overlapping community detection using seed set expansion. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, pages 2099–2108. ACM, 2013.
- [33] J. Xie, S. Kelley, and B. K. Szymanski. Overlapping community detection in networks: The state-of-the-art and comparative study. *Acm computing surveys (csur)*, 45(4):43, 2013.
- [34] J. Yang and J. Leskovec. Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 587–596. ACM, 2013.
- [35] J. Yang and J. Leskovec. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1):181–213, 2015.