

Clustering Large Attributed Graphs: An Efficient Incremental Approach

Yang Zhou, Hong Cheng, Jeffrey Xu Yu

Department of Systems Engineering and Engineering Management

The Chinese University of Hong Kong, Hong Kong, China

{zhou, hcheng, yu}@se.cuhk.edu.hk

Abstract—In recent years, many networks have become available for analysis, including social networks, sensor networks, biological networks, etc. Graph clustering has shown its effectiveness in analyzing and visualizing large networks. The goal of graph clustering is to partition vertices in a large graph into clusters based on various criteria such as vertex connectivity or neighborhood similarity. Many existing graph clustering methods mainly focus on the topological structures, but largely ignore the vertex properties which are often heterogeneous. Recently, a new graph clustering algorithm, *SA-Cluster*, has been proposed which combines structural and attribute similarities through a unified distance measure. *SA-Cluster* performs matrix multiplication to calculate the random walk distances between graph vertices. As the edge weights are iteratively adjusted to balance the importance between structural and attribute similarities, matrix multiplication is repeated in each iteration of the clustering process to recalculate the random walk distances which are affected by the edge weight update.

In order to improve the efficiency and scalability of *SA-Cluster*, in this paper, we propose an efficient algorithm *Inc-Cluster* to incrementally update the random walk distances given the edge weight increments. Complexity analysis is provided to estimate how much runtime cost *Inc-Cluster* can save. Experimental results demonstrate that *Inc-Cluster* achieves significant speedup over *SA-Cluster* on large graphs, while achieving exactly the same clustering quality in terms of intra-cluster structural cohesiveness and attribute value homogeneity.

Keywords—graph clustering; incremental computation;

I. INTRODUCTION

Graphs are popularly used to model structural relationship between objects in many application domains such as web, social networks, sensor networks, biological networks and communication networks, etc. Graph clustering has received a lot of attention recently with many proposed clustering algorithms [1], [2], [3], [4], [5]. Clustering on a large graph aims to partition the graph into several densely connected components. Typical applications of graph clustering include community detection in social networks, identification of functional modules in large protein-protein interaction networks, etc. Many existing graph clustering methods mainly focus on the topological structure of a graph so that each partition achieves a cohesive internal structure. Such methods include clustering based on normalized cuts [1], modularity [2], structural density [3] or flows [4]. On the other hand, a recent graph summarization method [6] aims to partition

a graph according to attribute similarity, so that nodes with the same attribute values are grouped into one partition.

In many real applications, both the graph topological structure and the vertex properties are important. For example, in a social network, vertex properties describe roles of a person while the topological structure represents relationships among a group of people. The graph clustering and summarization approaches mentioned above consider only one aspect of the graph properties but ignore the other. As a result, the clusters thus generated would either have a rather random distribution of vertex properties within clusters, or have a rather loose intra-cluster structure. An ideal graph clustering should generate clusters which have a *cohesive* intra-cluster structure with *homogeneous* vertex properties, by balancing the structural and attribute similarities.

Figure 1 shows an example of a coauthor graph where a vertex represents an author and an edge represents the coauthor relationship between two authors. In addition, there are an author ID, research topic and age associated with each author. The research topic and age are considered as attributes to describe the vertex properties. As we can see, authors r_1 – r_7 work on *XML*, authors r_9 – r_{11} work on *Skyline* and r_8 works on both. In addition, each author has a range value to describe his/her age. The problem studied in this paper is to cluster a graph associated with attributes (called an *attributed graph*), such as the example in Figure 1, based on both structural and attribute similarities. The goal is to partition the graph into k clusters with cohesive intra-cluster structures and homogeneous attribute values. The problem is quite challenging because structural and attribute similarities are two seemingly independent, or even conflicting goals – in our example, authors who collaborate with each other may have different values on research topics and age; while authors who work on the same topics or who are in a similar age may come from different groups with no collaborations. It is not straightforward to balance these two objectives.

In a recent work, Zhou et al. have proposed *SA-Cluster* [5], a graph clustering algorithm by combining structural and attribute similarities. A set of attribute vertices and attribute edges are added to the original graph. With such graph augmentation, the attribute similarity is transformed to vertex vicinity in the graph – two vertices which share an attribute value are connected by a common attribute




Figure 1. A Coauthor Network with Two Attributes “Topic” and “Age”

vertex. A neighborhood random walk model, which measures the vertex closeness on the augmented graph through both structure edges and attribute edges, unifies the two similarities. Then SA-Cluster uses the random walk distance as the vertex similarity measure and performs clustering by following the K-Medoids framework. As different attributes may have different degrees of importance, a weight ω_i , which is initialized to 1.0, is assigned to the attribute edges corresponding to attribute a_i . The attribute edge weights $\{\omega_1, \dots, \omega_m\}$ are updated in each iteration of the clustering process, to reflect the importance of different attributes. In the above example, after the first iteration, the weight of research topic will be increased to a larger value while the weight of age will be decreased, as research topic has better clustering tendency than age. Accordingly, the transition probabilities on the graph are affected iteratively with the attribute weight adjustments. Thus the random walk distance matrix needs to be *recalculated in each iteration of the clustering process*. Since the random walk distance calculation involves matrix multiplication, which has a time complexity of $O(n^3)$, the repeated random walk distance calculation causes a non-trivial computational cost in SA-Cluster. We find in the experiments that the random walk distance computation takes 98% of the total clustering time in SA-Cluster.

With a careful study of the weight self-adjustment mechanism in [5], we have observed that the weight increments only affect the attribute edges in the augmented graph, while the structure edges are not affected. Motivated by this, in this paper, we aim to improve the efficiency and scalability of SA-Cluster with a proposed efficient incremental computation algorithm Inc-Cluster to update the random walk distance matrix. The core idea is to compute the full random walk distance matrix only once at the beginning of the clustering process. Then in each following iteration of clustering, given the attribute weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$, we use Inc-Cluster to update the original random walk distance matrix, instead of re-calculating the matrix from scratch. This incremental computation problem is quite challenging. Existing incremental approaches [7], [8] cannot be directly applied to solve our problem, as they partition the graph into a changed part and an unchanged part. But in our problem it is hard to find such a clear boundary between the changed and the unchanged parts on the graph, because the effect of edge weight adjustments is

propagated widely to the whole graph in multiple steps. The distance between any pair of vertices may be affected. With the proposed Inc-Cluster algorithm, we can divide the graph clustering algorithm into two phases: an offline phase at the beginning of clustering for the full random walk distance matrix computation which is relatively expensive, and an online phase for the fast iterative clustering process with the incremental matrix calculation which is much cheaper. The main contributions of this paper are summarized below.

- 1) We study the problem of incremental computation of the random walk distance matrix in the context of graph clustering with structural and attribute similarities. We propose an efficient algorithm Inc-Cluster to incrementally update the random walk distance matrix given the attribute weight increments. By analyzing how the transition probabilities are affected by the weight increments, the random walk distance matrix is divided into submatrices for incremental update. Importantly, the incremental approach is also applicable to fast random walk computation in continuously evolving graphs with vertex/edge insertion and deletion.
- 2) Complexity analysis is provided to quantitatively estimate the upper bound and the lower bound of the number of elements in the random walk distance matrix that remain unchanged. The upper bound and lower bound correspond to the best case and the worst case of the incremental approach respectively.
- 3) We perform extensive evaluation of the incremental approach on real large graphs, demonstrating that our method Inc-Cluster is able to achieve significant speedup over SA-Cluster, while achieving exactly the same clustering quality in terms of intra-cluster structural cohesiveness and attribute value homogeneity.

The rest of the paper is organized as follows. We review related work on graph clustering in Section II. Section III introduces preliminary concepts and analyzes the runtime cost of SA-Cluster. Section IV presents our proposed incremental algorithm Inc-Cluster, followed by a complexity analysis in Section V. Section VI presents extensive experimental results. Finally, Section VII concludes the paper.

II. RELATED WORK

Many graph clustering techniques have been proposed which mainly focused on the topological structures based on various criteria including normalized cuts [1], modularity [2], structural density [3] or stochastic flows [4]. The clustering results contain densely connected components within clusters. However, such methods usually ignore vertex attributes in the clustering process. On the other hand, Tian et al. [6] proposed OLAP-style aggregation approaches to summarize large graphs by grouping nodes based on user-selected attributes. This method achieves homogeneous attribute values within clusters, but ignores the intra-cluster topological structures. Recently, Zhou et al. have proposed

a graph clustering algorithm, SA-Cluster [5], based on both structural and attribute similarities. Experimental results have shown that SA-Cluster achieves a good balance between structural cohesiveness and attribute homogeneity.

Other recent studies on graph clustering include the following. Sun et al. [9] proposed GraphScope which is able to discover communities in large and dynamic graphs, as well as to detect the changing time of communities. Sun et al. [10] proposed an algorithm, RankClus, which integrates clustering with ranking in large-scale information network analysis. Navlakha et al. [11] proposed a graph summarization method using the MDL principle.

The concept of random walk has been widely used to measure vertex distances. Jeh and Widom [12] designed a measure called SimRank, which defines the similarity between two vertices in a graph by their neighborhood similarity. Pons and Latapy [13] proposed to use short random walks of length l to measure the similarity between two vertices in a graph for community detection. Desikan et al. [7] proposed an incremental algorithm to compute PageRank for the evolving Web graph by partitioning the graph into a changed part and an unchanged part. [8] computes the local PageRank scores on a subgraph by assuming that the scores of external pages are known.

III. PRELIMINARY CONCEPTS

In this section, we first introduce the problem formulation of graph clustering considering both structural and attribute similarities. We then give a brief review of an earlier algorithm SA-Cluster by Zhou et al. [5] and analyze the computational cost. Our proposed approach to handle the computational bottleneck is outlined.

A. Attribute Augmented Graph

Definition 1 (Attributed Graph): An attributed graph is denoted as $G = (V, E, \Lambda)$, where V is the set of vertices, E is the set of edges, and $\Lambda = \{a_1, \dots, a_m\}$ is the set of m attributes associated with vertices in V for describing vertex properties. A vertex $v \in V$ is associated with an attribute vector $[a_1(v), \dots, a_m(v)]$ where $a_j(v)$ is the attribute value of vertex v on attribute a_j .

Attributed graph clustering is to partition an attributed graph G into k disjoint subgraphs $\{G_i = (V_i, E_i, \Lambda)\}_{i=1}^k$, where $V = \bigcup_{i=1}^k V_i$ and $V_i \cap V_j = \emptyset$ for any $i \neq j$. A desired clustering of an attributed graph should achieve a good balance between the following two objectives: (1) vertices within one cluster are close to each other in terms of structure, while vertices between clusters are distant from each other; and (2) vertices within one cluster have similar attribute values, while vertices between clusters could have quite different attribute values.

[5] proposed an *attribute augmented graph* to represent vertex associated attributes explicitly as *attribute vertices and edges*. In this paper we follow the same representation.




Figure 2. Attribute Augmented Graph

Definition 2 (Attribute Augmented Graph): Given an attributed graph $G = (V, E, \Lambda)$ with a set of attributes $\Lambda = \{a_1, \dots, a_m\}$. The domain of attribute a_i is $Dom(a_i) = \{a_{i1}, \dots, a_{in_i}\}$ with a size of $|Dom(a_i)| = n_i$. An attribute augmented graph is denoted as $G_a = (V \cup V_a, E \cup E_a)$ where $V_a = \{v_{ij}\}_{i=1, j=1}^{m, n_i}$ is the set of attribute vertices and $E_a \subseteq V \times V_a$ is the set of attribute edges. An attribute vertex $v_{ij} \in V_a$ represents that attribute a_i takes the j^{th} value. An attribute edge $(v_i, v_{jk}) \in E_a$ iff $a_j(v_i) = a_{jk}$, i.e., vertex v_i takes the value of a_{jk} on attribute a_j . Accordingly, a vertex $v \in V$ is called a structure vertex and an edge $(v_i, v_j) \in E$ is called a structure edge.

Figure 2 is an attribute augmented graph on the coauthor network example. Two attribute vertices v_{11} and v_{12} representing the topics “XML” and “Skyline” are added. Authors with corresponding topics are connected to the two vertices respectively in dashed lines. We omit the attribute vertices and edges corresponding to the age attribute, for the sake of clear presentation.

B. A Unified Random Walk Distance

In this paper we also use the neighborhood random walk model on the attribute augmented graph G_a to compute a unified distance between vertices in V . The random walk distance between two vertices $v_i, v_j \in V$ is based on the paths consisting of both structure and attribute edges. Thus it effectively combines the structural proximity and attribute similarity of two vertices into one unified measure. The transition probability matrix P_A on G_a is defined as follows.

A structure edge $(v_i, v_j) \in E$ is of a different type from an attribute edge $(v_i, v_{jk}) \in E_a$. The m attributes may also have different importance. Therefore, they may have different degree of contributions in random walk distance. Without loss of generality, we assume that a structure edge has a weight of ω_0 , attribute edges corresponding to a_1, a_2, \dots, a_m have an edge weight of $\omega_1, \omega_2, \dots, \omega_m$, respectively. Therefore, the transition probability from vertex v_i to vertex v_j through a structure edge is

$$p_{v_i, v_j} = \begin{cases} \frac{\omega_0}{|N(v_i)| * \omega_0 + \omega_1 + \dots + \omega_m}, & \text{if } (v_i, v_j) \in E \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where $N(v_i)$ represents the set of structure vertices connected to v_i . Similarly, the transition probability from v_i to

v_{jk} through an attribute edge is

$$p_{v_i, v_{jk}} = \begin{cases} \frac{\omega_j}{|N(v_i)| * \omega_0 + \omega_1 + \dots + \omega_m}, & \text{if } (v_i, v_{jk}) \in E_a \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

The transition probability from v_{ik} to v_j through an attribute edge is

$$p_{v_{ik}, v_j} = \begin{cases} \frac{1}{|N(v_{ik})|}, & \text{if } (v_{ik}, v_j) \in E_a \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

The transition probability between two attribute vertices v_{ip} and v_{jq} is 0 as there is no edge between attribute vertices.

$$p_{v_{ip}, v_{jq}} = 0, \forall v_{ip}, v_{jq} \in V_a \quad (4)$$

The transition probability matrix P_A is a $|V \cup V_a| \times |V \cup V_a|$ matrix, where the first $|V|$ rows (columns) correspond to the structure vertices and the rest $|V_a|$ rows (columns) correspond to the attribute vertices. For the ease of presentation, P_A is represented as

$$P_A = \begin{bmatrix} P_{V_1} & A_1 \\ B_1 & O \end{bmatrix} \quad (5)$$

where P_{V_1} is a $|V| \times |V|$ matrix representing the transition probabilities defined by Eq.(1); A_1 is a $|V| \times |V_a|$ matrix representing the transition probabilities defined by Eq.(2); B_1 is a $|V_a| \times |V|$ matrix representing the transition probabilities defined by Eq.(3); and O is a $|V_a| \times |V_a|$ zero matrix.

Definition 3 (Random Walk Distance Matrix): Let P_A be the transition probability matrix of an attribute augmented graph G_a . Given L as the length that a random walk can go, $c \in (0, 1)$ as the random walk restart probability, the unified neighborhood random walk distance matrix R_A is

$$R_A = \sum_{l=1}^L c(1-c)^l P_A^l \quad (6)$$

C. A Review of SA-Cluster

SA-Cluster adopts the *K-Medoids* clustering framework. After initializing the cluster centroids and calculating the random walk distance at the beginning of the clustering process, it repeats the following four steps until convergence.

- 1) Assign vertices to their closest centroids;
- 2) Update cluster centroids;
- 3) Adjust attribute edge weights $\{\omega_1, \dots, \omega_m\}$;
- 4) Re-calculate the random walk distance matrix R_A .

Different from traditional K-Medoids, SA-Cluster has two additional steps (i.e., steps 3-4): in each iteration, the attribute edge weights $\{\omega_1, \dots, \omega_m\}$ are automatically adjusted to reflect the clustering tendencies of different attributes. Interested readers can refer to [5] for the proposed mechanism for weight adjustment. According to Eq.(2), when the edge weights $\{\omega_1, \dots, \omega_m\}$ change, the transition

probability matrix P_A changes, so does the neighborhood random walk distance matrix R_A . As a result, the random walk distance matrix has to be re-calculated in each iteration due to the edge weight changes.

The cost analysis of SA-Cluster can be expressed as

$$t \cdot (T_{random_walk} + T_{centroid_update} + T_{assign})$$

where t is the number of iterations in the clustering process, T_{random_walk} is the cost of computing the random walk distance matrix R_A , $T_{centroid_update}$ is the cost of updating cluster centroids, and T_{assign} is the cost of assigning all points to cluster centroids.

For $T_{centroid_update}$ and T_{assign} the time complexity is $O(n)$ where $n = |V|$, since each of these two operations performs a linear scan of the graph vertices. On the other hand, the random walk distance calculation consists of matrix multiplications and additions, according to Eq.(6). Thus the time complexity for T_{random_walk} is $O(L \cdot n_a^3)$ where $n_a = |V \cup V_a|$ is the row (column) number of the transition probability matrix P_A . It is clear that T_{random_walk} is the dominant factor in the clustering process. The *repeated calculation* of random walk distance in each iteration can incur a non-trivial efficiency problem for SA-Cluster. We have observed that computing the random walk distance takes 98% of the total clustering time in SA-Cluster.

D. Our Solution: An Incremental Approach

The computational bottleneck in the random walk distance computation motivates us to seek alternative solutions with a lower cost. A natural direction to explore is “*can we avoid repeated calculation of random walk distance in the clustering process?*” The goal is to reduce the number of random walk distance calculation. We have observed that the attribute weight adjustments only change the transition probabilities of the attribute edges, but not those of the structure edges. This implies that many elements in the random walk distance matrix may remain unchanged. This property sheds light on the problem: we can design an *incremental calculation* approach to update the random walk distance matrix R_A iteratively. That is, given the original random walk distance matrix R_A and the weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$, efficiently calculate the increment matrix ΔR_A , and then get the updated random walk distance matrix $R_{N,A} = R_A + \Delta R_A$. In this process, we only calculate the *non-zero* elements in ΔR_A , i.e., those elements which are affected by the edge weight changes, but can ignore the unaffected parts of the original matrix. If the number of affected matrix elements is small, this incremental approach will be much more efficient than calculating the full matrix R_A from scratch in each iteration.

However, this incremental approach could be quite challenging, because the boundary between the changed part and the unchanged part of the graph is not clear. The attribute




Figure 3. Matrix Increment Series

weight adjustments will be propagated to the whole graph in L steps. Let us look at an example first.

Example 1: We select 1,000 authors from database, data mining, artificial intelligence and information retrieval with 3,782 edges for their collaborations. Each author has two attributes: “prolific” and “research topic”. The first attribute “prolific” contains two values of “highly prolific” and “low prolific”, and the second one “research topic” has 100 different values. Thus the augmented graph contains 1,000 structure vertices and 102 attribute vertices. The attribute edge weights for “prolific” and “research topic” are ω_1, ω_2 respectively. Figure 3 shows three matrices ΔP_A^1 , ΔP_A^2 and ΔP_A^{20} corresponding to the increments of the 1st, 2nd, and 20th power of the transition probability matrix, due to the attribute weight increments $\{\Delta\omega_1, \Delta\omega_2\}$. The blue dots represent non-zero elements and the red dashed lines divide each matrix into submatrices according to the block matrix representation in Eq.(5). As shown in Figure 3, ΔP_A^l becomes denser when l increases.

For ΔP_A^1 , the attribute weight increments only affect the transition probabilities in the submatrix A_1 , but cause no changes in the other three submatrices. Therefore, most elements in ΔP_A^1 are zero. ΔP_A^2 becomes denser with more non-zero elements. ΔP_A^{20} becomes even denser, which demonstrates that the effect of attribute weight increments is propagated to the whole graph through matrix multiplication.

Existing fast random walk [14] or incremental PageRank computation approaches [7], [8] can not be directly applied to our problem. Tong et al. [14] proposed an algorithm for fast random walk computation, which relies on partitioning the graph into k clusters apriori, to decompose the transition probability matrix into a within-partition one and a cross-partition one for a lower complexity. However, our graph clustering problem is much more difficult due to the augmented attribute edges and the iterative weight adjustments. The Incremental PageRank Algorithm (IPR) [7] computes PageRank for the evolving Web graph by partitioning the graph into a changed part and an unchanged part. The distribution of PageRank values in the unchanged part will not be affected. Two recent algorithms IdealRank and ApproxRank in [8] compute the PageRank scores in a subgraph, which is a small part of a global graph, by assuming that the scores of external pages, i.e., unchanged

pages, are known. Our incremental computation problem is much more challenging than the above problems. As we can see from Figure 3, although the edge weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$ affect a very small portion of the transition probability matrix P_A , (i.e., see ΔP_A^1), the changes are propagated widely to the whole graph through matrix multiplication (i.e., see ΔP_A^2 and ΔP_A^{20}). It is difficult to partition the graph into a changed part and an unchanged part and focus the computation on the changed part only.

IV. THE INCREMENTAL ALGORITHM

In this section, we will describe the incremental algorithm. According to Eq.(6), R_A is the weighted sum of a series of matrices P_A^l , where P_A^l is the l -th power of the transition probability matrix P_A , $l = 1, \dots, L$. Hence the problem of computing ΔR_A can be decomposed into the subproblems of computing ΔP_A^l for different l values. Therefore, our target is, given the original matrix P_A^l and the edge weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$, compute the increment ΔP_A^l .

A. Calculate 1st Power Matrix Increment ΔP_A^1

According to Eq.(5), the transition probability matrix P_A can be expressed as four submatrices P_{V_1} , A_1 , B_1 and O . Based on the transition probabilities defined in Eqs.(1)-(4) and the properties $\sum_{i=1}^m \omega_i = m$ and ω_0 is fixed, it is not hard to verify that the attribute weight increments only affect the transition probabilities in the submatrix A_1 , but cause no changes in the other three submatrices. Therefore, the increment of the transition probability matrix ΔP_A^1 is denoted as

$$\Delta P_A^1 = \begin{bmatrix} O & \Delta A_1 \\ O & O \end{bmatrix}$$

Consider a probability $p(v_i, v_{jk}) = \frac{\omega_j}{|N(v_i)| * \omega_0 + \omega_1 + \dots + \omega_m}$ as defined in Eq.(2). Given a new weight $\omega'_j = \omega_j + \Delta\omega_j$, the probability increment is

$$\Delta p(v_i, v_{jk}) = \frac{\Delta\omega_j}{|N(v_i)| * \omega_0 + \omega_1 + \dots + \omega_m} = \Delta\omega_j \cdot p(v_i, v_{jk}) \quad (7)$$

Eq.(7) holds because $\omega_j = 1.0$ and $\sum_{i=1}^m \omega_i = \sum_{i=1}^m \omega'_i = m$. Thus we denote $A_1 = [A_{a_1}, A_{a_2}, \dots, A_{a_m}]$ where A_{a_i} is a $|V| \times n_i$ matrix representing the transition probabilities from structure vertices in V to attribute vertices corresponding to attribute a_i . The column number n_i corresponds to

the n_i possible values in $Dom(a_i)$. An element $A_{a_i}(p, q)$ represents the transition probability from the p -th vertex $v_p \in V$ to the q -th value a_{iq} of a_i . Then ΔA_1 is equal to

$$\Delta A_1 = [\Delta\omega_1 \cdot A_{a_1}, \Delta\omega_2 \cdot A_{a_2}, \dots, \Delta\omega_m \cdot A_{a_m}] \quad (8)$$

where $\Delta\omega_i \cdot A_{a_i}$ is scalar multiplication, i.e., multiplying every element in A_{a_i} with $\Delta\omega_i$ according to Eq.(7). Then the new transition probability matrix $P_{N,A}$ after the edge weights change is represented as

$$P_{N,A} = \begin{bmatrix} P_{V_1} & A_1 + \Delta A_1 \\ B_1 & O \end{bmatrix} = \begin{bmatrix} P_{V_1} & A_{N,1} \\ B_1 & O \end{bmatrix}$$

B. Calculate l -th Power Matrix Increment ΔP_A^l

Similar to the computation of ΔP_A^1 , we can calculate ΔP_A^l ($l \geq 2$), with a more complicated computation. The original l -th power matrix $P_A^l = P_A^{l-1} \times P_A$ is represented as

$$\begin{aligned} P_A^l &= \begin{bmatrix} P_{V_{l-1}} & A_{l-1} \\ B_{l-1} & C_{l-1} \end{bmatrix} \times \begin{bmatrix} P_{V_1} & A_1 \\ B_1 & O \end{bmatrix} \\ &= \begin{bmatrix} P_{V_{l-1}}P_{V_1} + A_{l-1}B_1 & P_{V_{l-1}}A_1 \\ B_{l-1}P_{V_1} + C_{l-1}B_1 & B_{l-1}A_1 \end{bmatrix} \end{aligned}$$

Similarly, the new matrix $P_{N,A}^l = P_{N,A}^{l-1} \times P_{N,A}$ given the weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$ is

$$\begin{aligned} P_{N,A}^l &= \begin{bmatrix} P_{N,V_{l-1}} & A_{N,l-1} \\ B_{N,l-1} & C_{N,l-1} \end{bmatrix} \times \begin{bmatrix} P_{V_1} & A_{N,1} \\ B_1 & O \end{bmatrix} \\ &= \begin{bmatrix} P_{N,V_{l-1}}P_{V_1} + A_{N,l-1}B_1 & P_{N,V_{l-1}}A_{N,1} \\ B_{N,l-1}P_{V_1} + C_{N,l-1}B_1 & B_{N,l-1}A_{N,1} \end{bmatrix} \end{aligned}$$

Then the l -th power transition probability matrix increment ΔP_A^l is denoted as

$$\Delta P_A^l = \begin{bmatrix} \Delta P_{V_l} & \Delta A_l \\ \Delta B_l & \Delta C_l \end{bmatrix}$$

Based on the original matrix P_A^l and the new matrix $P_{N,A}^l$, the increment ΔP_{V_l} is

$$\begin{aligned} \Delta P_{V_l} &= (P_{N,V_{l-1}}P_{V_1} + A_{N,l-1}B_1) - (P_{V_{l-1}}P_{V_1} + A_{l-1}B_1) \\ &= (P_{V_{l-1}} + \Delta P_{V_{l-1}})P_{V_1} + (A_{l-1} + \Delta A_{l-1})B_1 \\ &\quad - (P_{V_{l-1}}P_{V_1} + A_{l-1}B_1) \\ &= \Delta P_{V_{l-1}}P_{V_1} + \Delta A_{l-1}B_1 \end{aligned}$$

The increment ΔB_l is

$$\begin{aligned} \Delta B_l &= (B_{N,l-1}P_{V_1} + C_{N,l-1}B_1) - (B_{l-1}P_{V_1} + C_{l-1}B_1) \\ &= (B_{l-1} + \Delta B_{l-1})P_{V_1} + (C_{l-1} + \Delta C_{l-1})B_1 \\ &\quad - (B_{l-1}P_{V_1} + C_{l-1}B_1) \\ &= \Delta B_{l-1}P_{V_1} + \Delta C_{l-1}B_1 \end{aligned}$$

The increment ΔA_l is

$$\begin{aligned} \Delta A_l &= P_{N,V_{l-1}}A_{N,1} - P_{V_{l-1}}A_1 \\ &= (P_{V_{l-1}} + \Delta P_{V_{l-1}})A_{N,1} - P_{V_{l-1}}A_1 \\ &= P_{V_{l-1}}\Delta A_1 + \Delta P_{V_{l-1}}A_{N,1} \end{aligned} \quad (9)$$

In Eq.(9), there is one component $P_{V_{l-1}}\Delta A_1$. As shown in Eq.(8), $\Delta A_1 = [\Delta\omega_1 \cdot A_{a_1}, \dots, \Delta\omega_m \cdot A_{a_m}]$, we then have

$$\begin{aligned} P_{V_{l-1}}\Delta A_1 &= P_{V_{l-1}}[\Delta\omega_1 \cdot A_{a_1}, \dots, \Delta\omega_m \cdot A_{a_m}] \\ &= [\Delta\omega_1 \cdot P_{V_{l-1}}A_{a_1}, \dots, \Delta\omega_m \cdot P_{V_{l-1}}A_{a_m}] \end{aligned}$$

Note that the submatrix A_l in P_A^l is computed by the following submatrix multiplication:

$$A_l = P_{V_{l-1}}A_1 + A_{l-1}O = P_{V_{l-1}}A_1$$

If we rewrite A_l as a series of $|V| \times n_i$ submatrices as $A_l = [A_{l,a_1}, A_{l,a_2}, \dots, A_{l,a_m}]$, then $A_{l,a_i} = P_{V_{l-1}}A_{a_i}$. As a result, $P_{V_{l-1}}\Delta A_1$ can be expressed as

$$\begin{aligned} P_{V_{l-1}}\Delta A_1 &= [\Delta\omega_1 \cdot P_{V_{l-1}}A_{a_1}, \dots, \Delta\omega_m \cdot P_{V_{l-1}}A_{a_m}] \\ &= [\Delta\omega_1 \cdot A_{l,a_1}, \dots, \Delta\omega_m \cdot A_{l,a_m}] \end{aligned}$$

Therefore, to compute $P_{V_{l-1}}\Delta A_1$ in Eq.(9), we only need to compute $[\Delta\omega_1 \cdot A_{l,a_1}, \dots, \Delta\omega_m \cdot A_{l,a_m}]$. The advantage is that $\Delta\omega_i \cdot A_{l,a_i}$ is scalar multiplication, which is much cheaper than the matrix multiplication on $P_{V_{l-1}}\Delta A_1$. Combining the above equations, we have

$$\Delta A_l = [\Delta\omega_1 \cdot A_{l,a_1}, \dots, \Delta\omega_m \cdot A_{l,a_m}] + \Delta P_{V_{l-1}}A_{N,1}$$

where the first part represents the attribute increment (i.e., the weight increments $\Delta\omega_i$'s on A_l), while the second part represents the accumulative increment from $\Delta P_{V_{l-1}}$.

Similarly, the increment ΔC_l is

$$\begin{aligned} \Delta C_l &= B_{N,l-1}A_{N,1} - B_{l-1}A_1 \\ &= (B_{l-1} + \Delta B_{l-1})A_{N,1} - B_{l-1}A_1 \\ &= B_{l-1}\Delta A_1 + \Delta B_{l-1}A_{N,1} \\ &= [\Delta\omega_1 \cdot C_{l,a_1}, \dots, \Delta\omega_m \cdot C_{l,a_m}] + \Delta B_{l-1}A_{N,1} \end{aligned}$$

where we represent $C_l = [C_{l,a_1}, C_{l,a_2}, \dots, C_{l,a_m}]$.

In summary, the l -th power matrix increment ΔP_A^l can be calculated based on: (1) the original transition probability matrix P_A and increment matrix ΔA_1 , (2) the $(l-1)$ -th power matrix increment ΔP_A^{l-1} , and (3) the original l -th power submatrices A_l and C_l . The key is that, if ΔA_1 and ΔP_A^{l-1} contain many zero elements, we can apply sparse matrix representation to speed up the matrix multiplication.

C. The Incremental Algorithm

Algorithm 1 presents the incremental algorithm for calculating the new random walk distance matrix $R_{N,A}$ given the original R_A and the weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$. The algorithm iteratively computes the increments ΔP_A^l for $l = 1, \dots, L$, and accumulates them into the increment matrix ΔR_A according to Eq.(6). Finally the new random walk distance matrix $R_{N,A} = R_A + \Delta R_A$ is returned.

The total runtime cost of the clustering process with Inc-Cluster can be expressed as

$$T_{random_walk} + (t-1) \cdot T_{inc} + t \cdot (T_{centroid_update} + T_{assign})$$

Algorithm 1 The Incremental Algorithm Inc-Cluster

Input: The original matrices $R_A, P_A, A_l, C_l, l = 2, \dots, L$
the attribute edge weight increments $\{\Delta\omega_1, \dots, \Delta\omega_m\}$
Output: The new random walk distance matrix $R_{N,A}$

```
1: Calculate  $\Delta P_A^1$  according to Eq.(8);
2:  $\Delta R_A = c(1-c)\Delta P_A^1$ ;
3: for  $l = 2, \dots, L$ 
4:    $\Delta P_{V_l} = \Delta P_{V_{l-1}} P_{V_l} + \Delta A_{l-1} B_l$ ;
5:    $\Delta B_l = \Delta B_{l-1} P_{V_l} + \Delta C_{l-1} B_l$ ;
6:    $\Delta A_l = [\Delta\omega_1 \cdot A_{l,a_1}, \dots, \Delta\omega_m \cdot A_{l,a_m}] + \Delta P_{V_{l-1}} A_{N,1}$ ;
7:    $\Delta C_l = [\Delta\omega_1 \cdot C_{l,a_1}, \dots, \Delta\omega_m \cdot C_{l,a_m}] + \Delta B_{l-1} A_{N,1}$ ;
8:    $\Delta R_{A+} = c(1-c)^l \Delta P_A^l$ ;
9: end for
10:return  $R_{N,A} = R_A + \Delta R_A$ ;
```

where T_{inc} is the time for incremental computation and T_{random_walk} is the time for computing the random walk distance matrix at the beginning of clustering. The speedup ratio r between SA-Cluster and Inc-Cluster is

$$\frac{t(T_{random_walk} + T_{centroid_update} + T_{assign})}{T_{random_walk} + (t-1)T_{inc} + t(T_{centroid_update} + T_{assign})}$$

Since $T_{inc}, T_{centroid_update}, T_{assign} \ll T_{random_walk}$, the speedup ratio is approximately

$$r \approx \frac{t \cdot T_{random_walk}}{T_{random_walk}} = t$$

Therefore, Inc-Cluster can improve the runtime cost of SA-Cluster by approximately t times, where t is the number of iterations in clustering.

V. COMPLEXITY ANALYSIS

In this section, we will perform some complexity analysis to estimate the number of zero elements in ΔP_A^l , as an indication to show how much cost Inc-Cluster can save. Intuitively, the more zero elements in the matrix increment ΔP_A^l , the less cost the incremental algorithm has. It is hard to give a closed form analytical result for a general $l \in \{1, \dots, L\}$, because we need to consider all possible length- l paths between any two vertices. So we focus on the analysis on ΔP_A^2 . Given a general attributed graph, we will provide an upper bound and a lower bound of the number of zero elements in ΔP_A^2 . This quantity directly affects the computational complexity of the incremental calculation.

Although we cannot provide the theoretical bounds for ΔP_A^l ($l > 2$), we observe in Figure 3 that the number of non-zero elements increases as l increases. However, we also observe from experiments, a large number of entries in ΔP_A^l approach to zero quickly when l increases, due to the multiplication of probabilities on the sequence of edges. Confirmed by our testing, over 75% of entries in ΔP_A^l become smaller than a very small threshold and can

be treated as zero. Therefore, practically the number of non-zero elements in ΔP_A^l is very small even for large l values.

In our analysis, we use the following notations: the m attributes a_1, \dots, a_m contain n_1, \dots, n_m values respectively. The number of structure vertices is $|V| = n$. Note that the following derived bounds do not make any assumption about the type of data or the value of m . Due to space limit, the detailed proofs of Lemmas 1 and 2 are omitted.

A. Upper Bound of The Number of Zero Elements in ΔP_A^2

Lemma 1: There are totally $\prod_{i=1}^m n_i$ combinations of attribute values among the m attributes, since an attribute a_i takes n_i values. Assume each combination has at least one vertex (without this assumption, we can find a special case with a trivial upper bound). When all vertices are evenly distributed in the $\prod_{i=1}^m n_i$ combinations of attribute values, i.e., each combination has $\frac{n}{\prod_{i=1}^m n_i}$ vertices, it gives the upper bound of the number of zero elements in ΔP_A^2 .

Theorem 1: The upper bound of the number of zero elements in ΔP_A^2 is

$$\frac{n^2 \times \prod_{i=1}^m (n_i - 1)}{\prod_{i=1}^m n_i} \quad (10)$$

Proof. If two vertices $v_i, v_j \in V$ have no common values on any attributes, then $\Delta P_A^2(i, j) = \Delta P_A^2(j, i) = 0$. Given one combination of the attribute values, there are $\prod_{i=1}^m (n_i - 1)$ combinations which do not share any attribute values with this combination. Since all vertices are evenly distributed in the $\prod_{i=1}^m n_i$ combinations of attribute values, there are $\frac{n}{\prod_{i=1}^m n_i}$ vertices belonging to each combination. Therefore, for any vertex v_i , the total number of vertices which do not share any attribute values with v_i is $\frac{n \times \prod_{i=1}^m (n_i - 1)}{\prod_{i=1}^m n_i}$. Accordingly, there are $\frac{n \times \prod_{i=1}^m (n_i - 1)}{\prod_{i=1}^m n_i}$ zero elements in $\Delta P_A^2(i, \cdot)$. Since there are totally n vertices in the graph, the total number of zero elements in ΔP_A^2 is

$$\frac{n^2 \times \prod_{i=1}^m (n_i - 1)}{\prod_{i=1}^m n_i}$$

$\frac{n^2 \times \prod_{i=1}^m (n_i - 1)}{\prod_{i=1}^m n_i}$ is in the scale of $O(n^2)$, which implies that most elements in ΔP_A^2 do not change. This corresponds to the best case of the incremental computation, since only a small number of elements in ΔP_A^2 need to be updated.

B. Lower Bound of The Number of Zero Elements in ΔP_A^2

Lemma 2: Assume each attribute value combination has at least one vertex. Among the $\prod_{i=1}^m n_i$ combinations of attribute values, assume for each of the first $\prod_{i=1}^m n_i - 1$ combinations, there exists exactly one vertex with the attribute vector corresponding to that combination. The remaining $n - (\prod_{i=1}^m n_i - 1)$ vertices have the same attribute vector corresponding to the last combination. This case gives the lower bound of the number of zero elements in ΔP_A^2 .

Theorem 2: The lower bound of the number of zero elements in ΔP_A^2 is

$$(2n - \prod_{i=1}^m n_i) \times \prod_{i=1}^m (n_i - 1) \quad (11)$$

Proof. Without loss of generality, we assume that exactly one vertex belongs to each of the first $\prod_{i=1}^m n_i - 1$ combinations of the attribute values. The set of such vertices is denoted as S . The set of the remaining vertices belonging to the last combination of attribute values is denoted as T . Let $S = S_1 \cup S_2$ where S_1 is the set of vertices which do not share any attribute values with vertices in T ; S_2 is the set of vertices which share one or more attribute values with vertices in T .

There are three cases to be discussed in the following to count the number of zero elements in ΔP_A^2 .

Case 1. Consider two vertices u and v . If u, v do not share the value on an attribute a_i , then v can take any of the other $n_i - 1$ values except the value taken by u . Since vertices in T and S_1 do not share any attribute values on the m attributes, there are totally $\prod_{i=1}^m (n_i - 1)$ combinations of attribute values that do not share with vertices in T . As we have assumed that there is exactly one vertex for each of such combinations, the size of S_1 is $|S_1| = \prod_{i=1}^m (n_i - 1)$ and the size of T is $|T| = n - (\prod_{i=1}^m n_i - 1)$. If two vertices v_i, v_j have no common values on any attributes, then $\Delta P_A^2(i, j) = \Delta P_A^2(j, i) = 0$. Therefore, $\forall v_i \in T, \forall v_j \in S_1, \Delta P_A^2(i, j) = 0$ and $\Delta P_A^2(j, i) = 0$. The number of such elements between T and S_1 is

$$LB_1 = 2|T| \times |S_1| = 2(n - (\prod_{i=1}^m n_i - 1)) \times \prod_{i=1}^m (n_i - 1)$$

Case 2. There exist some vertices in S which do not share any attribute values with any vertex in S_1 . We denote this set as $S_0, S_0 \subset S$. The size of S_0 is $|S_0| = \prod_{i=1}^m (n_i - 1) - 1$. So the total number of zero elements in ΔP_A^2 is:

$$2|S_1| \times |S_0| = 2 \prod_{i=1}^m (n_i - 1) \times (\prod_{i=1}^m (n_i - 1) - 1)$$

Since $S_0 \cap S_1 \neq \emptyset$, the above number double counts the following case: $v_i, v_j \in S_1$ and v_i, v_j do not share any attribute values. As a result, we have to deduct from the above $\prod_{i=1}^m (n_i - 1) \times \prod_{i=1}^m (n_i - 2)$ elements. Finally the number of zero elements in ΔP_A^2 in case 2 is

$$LB_2 = 2 \prod_{i=1}^m (n_i - 1) \times (\prod_{i=1}^m (n_i - 1) - 1) - \prod_{i=1}^m (n_i - 1) \times \prod_{i=1}^m (n_i - 2)$$

Case 3. There exist some vertices in S which do not share any attribute values with those in S_2 . The size of S_2 is $|S_2| = \prod_{i=1}^m n_i - 1 - \prod_{i=1}^m (n_i - 1)$. So the total number of zero elements in ΔP_A^2 in case 3 is:

$$(\prod_{i=1}^m n_i - 1 - \prod_{i=1}^m (n_i - 1)) \times \prod_{i=1}^m (n_i - 1)$$

However, the elements between any $v_i \in S_1, v_j \in S_2$ have been counted in case 2. So we should deduct the repeated counts. For a vertex $v_i \in S_1$, there are $\prod_{i=1}^m (n_i - 1) - \prod_{i=1}^m (n_i - 2) - 1$ vertices in S_2 which do not share any attribute values with it. Thus the number of repeated counts is $(\prod_{i=1}^m (n_i - 1) - \prod_{i=1}^m (n_i - 2) - 1) \times \prod_{i=1}^m (n_i - 1)$. Finally the number of zero elements for case 3 is

$$LB_3 = (\prod_{i=1}^m n_i + \prod_{i=1}^m (n_i - 2) - 2 \prod_{i=1}^m (n_i - 1)) \times \prod_{i=1}^m (n_i - 1)$$

By adding up LB_1, LB_2 and LB_3 , we can generate the lower bound of the number of zero elements in ΔP_A^2 .

$$LB = LB_1 + LB_2 + LB_3 = (2n - \prod_{i=1}^m n_i) \times \prod_{i=1}^m (n_i - 1)$$

As m and $n_i, i = 1, \dots, m$, are usually much smaller than n , LB is in the scale of $O(n)$, which is $\ll n^2$, the number of elements in ΔP_A^2 . Thus the lower bound corresponds to the worst case of the incremental computation, since most elements in ΔP_A^2 need to be updated.

VI. EXPERIMENTAL STUDY

In this section, we performed extensive experiments to evaluate the performance of Inc-Cluster on real graph data. All experiments were done in Matlab on a Dell PowerEdge R900 server with 2.67GHz six-core CPU and 128GB main memory running Windows Server 2008.

A. Experimental Datasets

We use the DBLP Bibliography data with 10,000 authors from four research areas of database, data mining, information retrieval and artificial intelligence. We build a coauthor graph where nodes represent authors and edges represent their coauthor relationships. In addition, we use two relevant attributes: *prolific* and *primary topic*. For “prolific”, authors with ≥ 20 papers are labeled as highly prolific; authors with ≥ 10 and < 20 papers are labeled as prolific and authors with < 10 papers are labeled as low prolific. For “primary topic”, we use a topic modeling approach [15] to extract 100 topics from a document collection composed of paper titles from the selected authors. Each extracted topic consists of a probability distribution of keywords which are most representative of the topic. Then each author will have one out of 100 topics as his/her primary topic.

We also use a larger DBLP dataset with 84,170 authors, selected from the following areas: database, data mining, information retrieval, machine learning, artificial intelligence, computer systems, theory, computer vision, architecture, programming language, networking, simulation, natural language processing, multimedia, and human-computer interaction. The coauthor graph and the vertex attributes are defined similarly as in the 10,000 coauthor network.

B. Comparison Methods and Evaluation

We tested the following algorithms for the clustering quality and efficiency comparison.

- **Inc-Cluster** Our proposed algorithm which incrementally updates the random walk distance matrix.
- **SA-Cluster** The non-incremental graph clustering algorithm [5] which considers both structural and attribute similarities.
- **S-Cluster** The graph clustering algorithm which only considers topological structure. Random walk distance is used to measure vertex closeness while attribute similarity is ignored.
- **W-Cluster** A fictitious clustering algorithm which combines structural and attribute similarities through a weighted function as $\alpha \cdot d_S(v_i, v_j) + \beta \cdot d_A(v_i, v_j)$, where $d_S(v_i, v_j)$ is the random walk distance, and $d_A(v_i, v_j)$ is their attribute similarity, and the weighting factors are $\alpha = \beta = 0.5$.
- **k-SNAP** The k-SNAP algorithm [6] that groups vertices with the same attribute values into one cluster.

Evaluation Measures We use two measures of *density* and *entropy* to evaluate the quality of clusters $\{V_i\}_{i=1}^k$ generated by different methods. The definitions are as follows.

$$density(\{V_i\}_{i=1}^k) = \sum_{i=1}^k \frac{|\{(v_p, v_q) | v_p, v_q \in V_i, (v_p, v_q) \in E\}|}{|E|}$$

$$entropy(\{V_i\}_{i=1}^k) = \sum_{i=1}^m \frac{\omega_i}{\sum_{p=1}^m \omega_p} \sum_{j=1}^k \frac{|V_j|}{|V|} entropy(a_i, V_j)$$

where $entropy(a_i, V_j) = -\sum_{n=1}^{n_i} p_{ijn} \log_2 p_{ijn}$ and p_{ijn} is the percentage of vertices in cluster j which have value a_{in} on attribute a_i . $entropy(\{V_i\}_{i=1}^k)$ measures the weighted entropy from all attributes over k clusters.


Besides the clustering quality comparison, we also compare the runtime efficiency of these methods.

C. Clustering Quality Comparison


Since SA-Cluster and Inc-Cluster generate the same clustering results, their quality results are shown in the same column in Figures 4 and 5.

Figure 4 (a) shows the density on the DBLP graph with 10,000 authors by different methods. The density values by SA-Cluster and Inc-Cluster are around 0.51 – 0.60, which are slightly lower than those of S-Cluster. The density values by W-Cluster and k-SNAP are much lower, in the range of 0.15 – 0.18. This shows the clusters generated by W-Cluster and k-SNAP have a very loose intra-cluster structure.

Figure 4 (b) shows the entropy comparison on DBLP with 10,000 authors. S-Cluster has the highest entropy around 2.7 – 3.0, because it partitions a graph without considering vertex attributes. SA-Cluster and Inc-Cluster have a low entropy around 1.1 – 1.2. W-Cluster has an even lower entropy but also a very low density. This is because its distance



(a) density (b) entropy
Figure 4. Cluster Quality on DBLP 10,000 Authors



(a) density (b) entropy
Figure 5. Cluster Quality on DBLP 84,170 Authors

function combines (or compromises) both structural and attribute similarities through a weighted function. However, as it is not clear how to set or tune the weighting factors α and β , it is hard to achieve an optimal result on W-Cluster. Since k-SNAP strictly enforces the attribute homogeneity in each cluster, k-SNAP achieves an entropy of 0.

Figures 5 (a) and (b) show the density and entropy on DBLP with 84,170 authors when $k = 400, 800, 1200, 1600$. These two figures have a similar trend with Figures 4 (a)–(b). SA-Cluster and Inc-Cluster achieve similar high density values (>0.90) with S-Cluster, but with much lower entropy. W-Cluster and k-SNAP achieve very low entropy (the entropy by k-SNAP is 0), but with very low density values at 0.2 – 0.3. The comparison on both density and entropy demonstrates that both SA-Cluster and Inc-Cluster achieve a very good balance between the structural cohesiveness and attribute homogeneity.

D. Clustering Efficiency Comparison

In this experiment, we compare the efficiency of different clustering algorithms. Figures 6 (a) and (b) show the clustering time on DBLP with 10,000 and 84,170 authors respectively. We make the following observations on the runtime costs of different methods. First, SA-Cluster is usually 3.1 – 3.8 times slower than Inc-Cluster, as it iteratively computes the random walk distance matrix from scratch. According to our analysis, the speedup ratio r is determined by the number of iterations in clustering, which are 3–6 iterations in the experiments. We have also observed that the random walk distance matrix computation takes 98% of the total clustering time in SA-Cluster. Second, S-Cluster and W-Cluster are usually faster than Inc-Cluster, as they compute the random walk distance only once on a smaller scale matrix (i.e., without augmentation with attribute vertices and edges). Third, the runtime of k-SNAP increases dramatically with k .

The statistics on the number of zero elements in ΔP_A^2




Figure 6. Clustering Efficiency




Figure 7. Runtime Comparison

also testify our previously proved bounds. On DBLP 10,000 data, there are 24M zero entries in ΔP_A^2 , while the theoretical upper and lower bounds are 66M and 4M, respectively. On DBLP 84,170 data, there are 4.4B zero entries, while the upper and lower bounds are 4.7B and 33M, respectively.

Figure 7 compares T_{random_walk} and T_{inc} to compute the random walk distance matrix R_A on DBLP with 84,170 authors. To illustrate the difference between the two approaches, we use a threshold δ to prune small values in P_A^l , i.e., $P_A^l(i, j) := 0$ if $P_A^l(i, j) \leq \delta$. Specifically, we set $\delta = 0.0001/x^l$ to progressively prune small values in P_A^l , $l \in \{1, \dots, L\}$. x is a decay factor because values in P_A^l become smaller and smaller as l increases. We set $x = 2.8-3.3$ as shown in the x -axis. The smaller x is, the more elements are set to 0 in P_A^l , thus the faster the matrix multiplication is. As shown in Figure 7, T_{random_walk} is very sensitive to the decay factor – the runtime increases dramatically with x , because more and more non-zero elements exist in P_A^l when x increases. On the other hand, T_{inc} remains stable as x increases, because T_{inc} is determined only by the number of zero elements in ΔP_A^l , but not by the number of zero elements in P_A^l . In other words, although many elements are non-zero in P_A^l when x is large, as long as most of them remain unchanged with the attribute weight update, there is little overhead for the incremental approach.

VII. CONCLUSION

In this paper, we propose an incremental algorithm Inc-Cluster to quickly compute a random walk distance matrix, in the context of graph clustering considering both structural and attribute similarities. To avoid recalculating the random walk distances from scratch in each iteration due to the attribute weight changes, we divide the transition probability matrix into submatrices and incrementally update each one. Time complexity analysis is provided to show the properties of Inc-Cluster. Experimental results show that Inc-Cluster

achieves significant speedup over SA-Cluster, while achieving the same clustering quality.

ACKNOWLEDGMENT

The work was supported in part by grants of the Research Grants Council of the Hong Kong SAR, China No. 419008 and the Chinese University of Hong Kong Direct Grants No. 2050446 and No. 2050473.

REFERENCES

- [1] J. Shi and J. Malik, “Normalized cuts and image segmentation,” in *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.
- [2] M. E. J. Newman and M. Girvan, “Finding and evaluating community structure in networks,” in *Phys. Rev. E* 69, 026113, 2004.
- [3] X. Xu, N. Yuruk, Z. Feng, and T. A. J. Schweiger, “Scan: a structural clustering algorithm for networks,” in *KDD*, 2007, pp. 824–833.
- [4] V. Satuluri and S. Parthasarathy, “Scalable graph clustering using stochastic flows: Applications to community discovery,” in *KDD*, 2009, pp. 737–745.
- [5] Y. Zhou, H. Cheng, and J. X. Yu, “Graph clustering based on structural/attribute similarities,” in *VLDB*, 2009, pp. 718–729.
- [6] Y. Tian, R. A. Hankins, and J. M. Patel, “Efficient aggregation for graph summarization,” in *SIGMOD*, 2008, pp. 567–580.
- [7] P. Desikan, N. Pathak, J. Srivastava, and V. Kumar, “Incremental page rank computation on evolving graphs,” in *WWW*, 2005, pp. 1094–1095.
- [8] Y. Wu and L. Raschid, “Approxrank: Estimating rank for a subgraph,” in *ICDE*, 2009, pp. 54–65.
- [9] J. Sun, C. Faloutsos, S. Papadimitriou, and P. S. Yu, “Graphscope: parameter-free mining of large time-evolving graphs,” in *KDD*, 2007, pp. 687–696.
- [10] Y. Sun, J. Han, P. Zhao, Z. Yin, H. Cheng, and T. Wu, “Rankclus: Integrating clustering with ranking for heterogeneous information network analysis,” in *EDBT*, 2009, pp. 565–576.
- [11] S. Navlakha, R. Rastogi, and N. Shrivastava, “Graph summarization with bounded error,” in *SIGMOD*, 2008, pp. 419–432.
- [12] G. Jeh and J. Widom, “SimRank: a measure of structural-context similarity,” in *KDD*, 2002, pp. 538–543.
- [13] P. Pons and M. Latapy, “Computing communities in large networks using random walks,” *J. Graph Algorithms and Applications*, vol. 10, no. 2, pp. 191–218, 2006.
- [14] H. Tong, C. Faloutsos, and J.-Y. Pan, “Fast random walk with restart and its applications,” in *ICDM*, 2006, pp. 613–622.
- [15] T. Hofmann, “Probabilistic latent semantic indexing,” in *SI-GIR*, 1998, pp. 50–57.