



Efficient Algorithms for Density Decomposition on Large Static and Dynamic Graphs

Yalong Zhang
Beijing Institute of Technology
Beijing, China
yalong-zhang@qq.com

Rong-Hua Li
Beijing Institute of Technology
Beijing, China
lironghuabit@126.com

Qi Zhang
Beijing Institute of Technology
Beijing, China
qizhangcs@bit.edu.cn

Hongchao Qin
Beijing Institute of Technology
Beijing, China
qhc.neu@gmail.com

Guoren Wang
Beijing Institute of Technology
Beijing, China
wanggrbit@126.com

ABSTRACT

Locally-densest subgraph (LDS) decomposition is a fundamental decomposition in graph analysis that finds numerous applications in various domains, including community detection, fraud detection, graph querying, and graph visualization. However, the LDS decomposition is computationally challenging for both static and dynamic graphs. Furthermore, the LDS decomposition often produces an excessive number of dense subgraph layers, leading to the unnecessary partition of tightly-connected subgraphs. To address these limitations, an alternative concept called density decomposition was proposed, which can generate a more reasonable number of dense subgraph layers. However, the state-of-the-art algorithm for density decomposition requires $O(m^2)$ time (m is the number of edges of the graph), which is very costly for large graphs. In this paper, we conduct an in-depth investigation of density decomposition and propose efficient algorithms for computing it on both static and dynamic graphs. First, we establish a novel relationship between density decomposition and LDS decomposition. Second, based on these relationships, we propose novel algorithms to compute the density decomposition on static graphs with carefully designed network flow and divide-and-conquer techniques. Our proposed static algorithms significantly reduce the time complexity to $O(m^{3/2} \log p)$ (p is often a very small constant in real-world graphs). Third, for dynamic graphs, we develop three dynamic algorithms with efficient $O(m)$ time complexity. Extensive experiments on several large real-world graphs demonstrate the high efficiency, scalability, and effectiveness of the proposed algorithms.

PVLDB Reference Format:

Yalong Zhang, Rong-Hua Li, Qi Zhang, Hongchao Qin, and Guoren Wang. Efficient Algorithms for Density Decomposition on Large Static and Dynamic Graphs. PVLDB, 17(11): 2933 - 2945, 2024.
doi:10.14778/3681954.3681974

PVLDB Artifact Availability:

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

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. 17, No. 11 ISSN 2150-8097.
doi:10.14778/3681954.3681974

1 INTRODUCTION

Real-world graphs are typically overall sparse but contain small dense subgraphs. Discovering these dense subgraphs has a wide range of applications, including community detection [1, 16, 26, 34, 45, 59], fraud detection [5], graph querying [18, 33], and graph visualization [2, 60]. Various models have been introduced to characterize dense subgraphs, such as the clique [24], k -core [3, 36, 40, 46], k -truss [19, 32], locally densest subgraph [37, 43, 54], k -edge connected subgraph [12, 13], k -plex [21, 50], k -defective clique [20], k -club [41], and so on.

Among them, k -core perhaps stands as the most popular dense graph model due to its concise definition and linear time complexity. However, the k -core decomposition, defined solely based on degrees, falls short of effectively capturing the density structure of a graph, where density is measured as the ratio of the number of edges to the number of vertices. In addition, graph decomposition should reasonably partition the graph into communities based on density. However, it has been demonstrated that k -core decomposition does not fulfill this requirement. For example, as shown in Fig. 1a, the graph comprises two densely-connected communities. However, community #1 and community #2 are partitioned by C_4 and C_3 , respectively, which constitutes an unreasonable partitioning approach. Besides, the top layer of core decomposition, C_4 , contains 10 edges and 5 vertices, yielding a density of 10/5. Nevertheless, the subgraph exhibiting the highest density in this graph is the whole community #1 with a density of 19/8. Clearly, the core decomposition is not density-based, and thus may not accurately capture the dense communities.

To address this limitation, Tatti [54] introduced a concept of *locally-densest subgraph (LDS) decomposition*, dividing the graph into a set of LDSes with hierarchical features, i.e., a smaller subgraph must be contained in a larger subgraph, thereby forming a decomposition. It is density-based and ensures that the top layer is exactly the densest subgraph in the entire graph. However, LDS decomposition has three notable drawbacks: (1) The worst-case time complexity for computing it is as high as $O(n^2m)$, where n and m denote the number of vertices and edges, which is costly for large graphs; (2) The number of decomposed LDSes may be excessive and can only be bounded by $O(n)$, typically leading to unnecessary partition of densely-connected communities. For example, in Fig. 1b, B_1 separates the tightly connected subgraph $\{v_9, \dots, v_{16}\}$, which is not necessary; (3) It is very challenging to exactly maintain the LDSes on dynamic graphs due to the strictness of its definition.

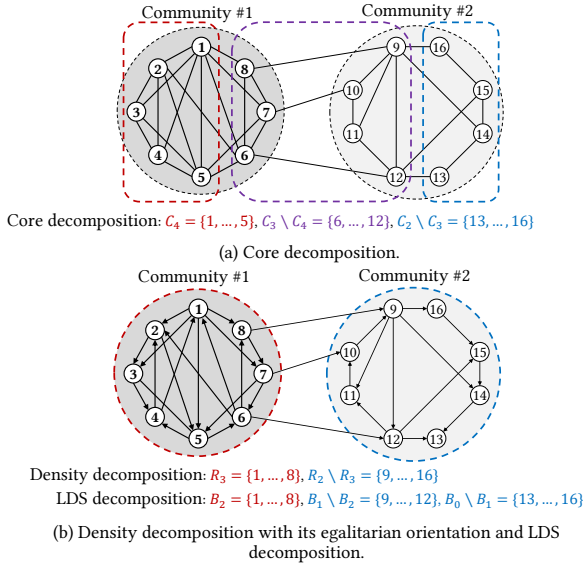


Figure 1: Example graph $G = (V, E)$ and its decompositions, where density decomposition exhibits the best performance.

On the other hand, Borradaile et al. [11] introduced an alternative density-based decomposition, known as *density decomposition*, which partitions the vertices based on their indegree and connectivity in a so-called egalitarian orientation [10]. In [11], Borradaile et al. proposed an $O(m^2)$ time algorithm to compute the density decomposition, which is clearly costly for large graphs. Moreover, Borradaile et al. did not establish a clear relationship between the density decomposition and other dense subgraph models.

Motivated by this, we for the first time conduct a thorough investigation of density decomposition, providing insights and a comprehensive understanding of its properties and applications. Specifically, we first establish a novel and interesting connection between density decomposition and LDS decomposition, which enables us to use the existing highly-optimized LDS decomposition algorithm to compute the density decomposition [37]. Based on these connections, we propose novel algorithms with network flow techniques to efficiently compute the density decomposition. Subsequently, we also develop novel algorithms with linear time complexity to efficiently maintain the density decomposition on dynamic graphs. Finally, we conduct extensive experiments to evaluate our algorithms using several large real-life graphs, and the results demonstrate the efficiency, scalability, and effectiveness of the proposed algorithms. In summary, the main contributions of this paper are as follows.

New theoretical results on density decomposition. We show that each layer of the density decomposition is exactly an LDS, making it part of the LDS decomposition. This crucial relationship reveals the fundamental similarity between these two decompositions. On the other hand, their key distinction lies in that density decomposition, through a round-up relationship, can naturally group LDSes with similar densities. This grouping significantly reduces the number of subgraphs, addressing the issue of unnecessary layers in the LDS decomposition. In particular, we show that the number of non-trivial layers generated by the density decomposition is equal to the *pseudoarboricity* of a graph, denoted as p , typically a small number. In addition, we also reveal an interesting

approximation relationship between density decomposition and k -core decomposition, indicating that core decomposition can be considered as a 2-approximation of the density decomposition.

Novel algorithms for computing density decomposition. We first propose a novel algorithm, called Flow, to compute the density decomposition, based on a carefully-designed network flow technique. Compared to the state-of-the-art $O(m^2)$ -time algorithm, Flow achieves a superior time complexity of $O(p \cdot m^{3/2})$. To further improve the efficiency, we propose a powerful divide-and-conquer technique and develop a new network flow-based algorithm, called Flow++. We show that such a new algorithm further reduces the time complexity to $O(m^{3/2} \log p)$.

Novel algorithms for maintaining density decomposition. We for the first time study the problem of maintaining the density decomposition on dynamic graphs. We discover a density decomposition update theorem, revealing that the insertion or deletion of an edge requires analyzing changes in just one layer of the density decomposition. Based on this, we develop an insertion algorithm Insert to handle edge insertions with a worst-case time complexity of $O(m)$. For edge deletion, we develop the Delete algorithm along with its improved version, Delete++, which incorporates several carefully designed pruning strategies. Both of these algorithms have $O(m)$ time complexity, thus they are very efficient for handling large dynamic graphs.

Extensive experiments. We conduct comprehensive experiments to evaluate the proposed algorithms using 8 real-world datasets. The results are summarized as follows: (1) Both the Flow and Flow++ algorithms are substantially faster than the state-of-the-art algorithm by at least two orders of magnitude for density composition. Moreover, even when compared to the highly-optimized LDS decomposition algorithm [37], our Flow++ can still achieve a remarkable speed improvement of one order of magnitude; (2) For density decomposition maintenance, both our insertion and deletion algorithms outperform the baselines by at least 5 orders of magnitude. We also conduct various case studies to demonstrate the effectiveness of density decomposition. The findings reveal that, unlike density-based methods, core decomposition often fails to identify densely-connected communities, leading to unsatisfactory performance in practical applications. LDS decomposition, with its excessive layers, tends to forcibly separate densely-connected communities and divide sparse areas unnecessarily. In contrast, density decomposition accurately captures the density structure of graphs, thereby effectively locating dense communities.

2 PRELIMINARIES

Consider an undirected graph $G = (V, E)$, where V and E are the vertex set and edge set respectively. Let $n = |V|$ and $m = |E|$ be the number of vertices and edges in G . An *orientation* of G can be obtained by assigning a direction to every edge in E , resulting in a directed graph $\vec{G} = (V, \vec{E})$. For example, the directed graph in Fig. 1b is an orientation of the undirected graph in Fig. 1a. To eliminate confusion, we use angle brackets $\langle u, v \rangle$ to represent a directed edge in \vec{G} , while round brackets (u, v) denote an undirected edge in G . For a directed (resp., undirected) graph G and a vertex u , we use $d_u(G)$ (d_u for brevity) to denote the indegree (resp., degree) of u . A *path* in a directed graph $\vec{G} = (V, \vec{E})$ is a sequence of vertices $s = v_0, v_1, \dots, v_{l-1}, t = v_l$, where $\langle v_{i-1}, v_i \rangle \in \vec{E}$, for $i = 1, \dots, l$. We denote such a path as $s \rightsquigarrow t$, and the length of $s \rightsquigarrow t$ is l . If there

exists a path $s \rightsquigarrow t$, we say that s can *reach* t . Besides, if an edge $\langle u, v \rangle$ is *reversed*, then it becomes $\langle v, u \rangle$. When a path is *reversed*, all edges in the path undergo a reversal. Before introducing the density decomposition, we first give the definitions of *reversible path* and *egalitarian orientation* as follows.

DEFINITION 1. (Reversible path) *Given a graph G and its orientation \vec{G} , for a path $s \rightsquigarrow t$ in \vec{G} , if $d_t(\vec{G}) - d_s(\vec{G}) \geq 2$, we call that the path $s \rightsquigarrow t$ is a reversible path.*

DEFINITION 2. (Egalitarian orientation) *Given a graph G and its orientation \vec{G} , if there is no reversible path in \vec{G} , then \vec{G} is an egalitarian orientation.*

Intuitively, an egalitarian orientation distributes the indegree of vertices in the most equitable manner, i.e., minimizing the indegree difference between vertices as much as possible. Note that if reverse a reversible path $s \rightsquigarrow t$, d_s increases by 1, d_t decreases by 1, and the indegree of other vertices does not change, making the indegree difference between s and t reduced by 2. When no reversible path exists, the indegree difference can not be reduced anymore. Based on the egalitarian orientation, we give the definition of density decomposition originally proposed by Borradaile et al. [11].

DEFINITION 3. (Density decomposition) [11] *Given an undirected graph $G = (V, E)$ and its arbitrary egalitarian orientation $\vec{G} = (V, \vec{E})$, the density decomposition of G is a set of subgraphs, denoted by $\mathcal{R} = \{R_k\}$, where for any non-negative integer k , $R_k \triangleq \{u \in V | d_u(\vec{G}) \geq k \text{ or } u \text{ can reach a vertex } v \text{ with } d_v(\vec{G}) \geq k\}$.*

To provide an intuition for Definition 3, we introduce the Theorem 1. This theorem specifies that any subgraph $S \subseteq R_k$ possesses a higher edge count, rendering it denser, while any subgraph T outside of R_k contains fewer edges, making it sparser. This intuitive analysis demonstrates that the density decomposition, defined by indegree and reachability, is inherently density-based. Due to space limitations, we have included the full proof in the full version of the paper, which can be found at https://github.com/Flydragonet/Density_Decomposition_Computation.

THEOREM 1. *Given a layer R_k of density decomposition, it satisfies: (i) For any $S \subseteq R_k$, the removal of S from R_k results in a deletion of at least $(k-1) \cdot |S|$ edges; (ii) For any $T \subseteq V \setminus R_k$, the inclusion of T into R_k leads to an increase of at most $(k-1) \cdot |T|$ edges.*

PROOF SKETCH. The proof uses the notations of E_\times and E_Δ , with definitions provided in Section 3. We aim to establish that $|E_\Delta(S, R_k \setminus S)| > (k-1) \cdot |S|$ and $(k-1) \cdot |T| \geq |E_\Delta(T, R_k)|$. Suppose \vec{G} is an egalitarian orientation of G . By Lemma 1, in \vec{G} , any edge in $E_\times(R_k, V \setminus R_k)$ must point towards $V \setminus R_k$. Therefore, $|E_\Delta(S, R_k \setminus S)| \geq \sum_{u \in S} d_u(\vec{G}) \geq (k-1) \cdot |S|$ and $(k-1) \cdot |T| \geq \sum_{u \in T} d_u(\vec{G}) \geq |E_\Delta(T, R_k)|$. It is critical to note that the first two inequalities cannot hold simultaneously; otherwise, all vertices in S would have an indegree of $k-1$, thereby cannot reach vertices with an indegree of at least k , contradicting the fact that $S \subseteq R_k$. Hence, $|E_\Delta(S, R_k \setminus S)| > (k-1) \cdot |S|$ and $(k-1) \cdot |T| \geq |E_\Delta(T, R_k)|$. \square

Based on Definition 3, we define *integral dense number* as follows.

DEFINITION 4. (Integral dense number, IDN) *For a vertex $u \in R_k \setminus R_{k+1}$, the IDN of u is defined as $\bar{r}_u = k$.*

Table 1: Frequently used notations.

Notation	Definition
$d_u(G), d_u(\vec{G})$	the degree of u in G , the indegree of u in \vec{G}
$s \rightsquigarrow t$	a path from s to t
r_u, \bar{r}_u	fractional dense number (FDN) of u , integral dense number (IDN) of u
$p(G)$ or p	the maximum integer k such that $R_k \neq \emptyset$, equaling to pseudoarboricity
$\rho(X)$	the density of X : $\rho(X) = E(X) / X $
$E_\times(X, Y)$	the cross edges between X and Y : if $X \cap Y = \emptyset$, $E_\times(X, Y) = \{(x, y) \in E x \in X, y \in Y\}$, otherwise $E_\times(X, Y) = E_\times(X, Y \setminus X)$
$E_\Delta(X, Y)$	the additional edges from X with respect to Y : if $X \cap Y = \emptyset$, $E_\Delta(X, Y) = E(X) \cup E_\times(X, Y)$, otherwise $E_\Delta(X, Y) = E_\Delta(X, Y \setminus X)$
$\rho(X, Y)$	the outer density of X with respect to Y : if $X \cap Y = \emptyset$, $\rho(X, Y) = E_\Delta(X, Y) / X $, otherwise $\rho(X, Y) = \rho(X, Y \setminus X)$

EXAMPLE 1. *Taking the egalitarian orientation shown in Fig. 1b as an example, only v_2 and v_5 have the indegree of 3. Thus, R_3 contains v_2, v_5 , and the vertices that can reach them, i.e., $R_3 = \{v_1, \dots, v_8\}$. For vertices v_1, \dots, v_{15} , each has an indegree of at least 2, and v_{16} can reach them, thus R_2 contains all vertices. Besides, R_1 and R_0 also contain all vertices, and for any integer $k > 3$, $R_k = \emptyset$. For Theorem 1, we consider any subgraph S of R_3 , such as $S = \{6, 7, 8\}$. Removing S from R_3 results in a loss of 9 edges from $E(R_3)$, constituting up to half of the total edge number $|E(R_3)|$. This indicates that R_3 is relatively dense and compact, as each of its subgraphs contains a substantial number of edges. Conversely, we consider any subgraph T outside of R_3 , for instance $T = \{9, 10, 11\}$. Incorporating T into R_3 would only increase the edge count in $E(R_3)$ by 5. This modest increment illustrates the relative sparsity of T and justifies its exclusion from R_3 .*

The density decomposition also possesses the following fundamental properties: (1) Given a graph, density decomposition is unique, i.e., a graph has exactly one density decomposition [11]; (2) It exhibits a hierarchical structure, i.e., for two non-negative integers $i > j$, R_i is a subset of R_j [11]. Furthermore, Borradaile et al. [11] also established the following lemma, proving some fundamental properties of the egalitarian orientation.

LEMMA 1. [11] *Given an undirected graph $G = (V, E)$ and its arbitrary egalitarian orientation $\vec{G} = (V, \vec{E})$, if $u \in R_k \setminus R_{k+1}$, then $d_u(\vec{G}) \in \{k, k-1\} = \{\bar{r}_u, \bar{r}_u - 1\}$. Additionally, if $(u, v) \in E$ and $\bar{r}_u < \bar{r}_v$, then in \vec{G} , (u, v) must be oriented as $\langle v, u \rangle$ instead of $\langle u, v \rangle$.*

For the number of layers of density decomposition, let p be the largest integer such that R_p is non-empty. According to [6, 56], we can derive that p equals the pseudoarboricity of G , whose definition is in [42]. As indicated in [7], the pseudoarboricity of real-world graphs is often very small. Therefore, the number of layers of the density decomposition, from $R_0 = V$ to $R_{p+1} = \emptyset$, is $p+2$, and the number of non-trivial layers, from R_1 to R_p , is p .

With these preliminaries, we formulate our problems as follows.

Problem definition. Given an undirected graph G , our goal is to efficiently compute the density decomposition of G , i.e., computing the IDNs for all vertex in G . For dynamic graphs with edge insertions or deletions, we aim to maintain the density decomposition, i.e., maintain the IDNs for all vertices.

3 NEW THEORETICAL INSIGHTS

In this section, we first establish an interesting connection between density decomposition and locally-densest subgraph (LDS) decomposition, allowing us to leverage existing highly-optimized LDS decomposition algorithms, such as [37] and [30], for computing density decomposition. Then, we will show a close connection between density decomposition and core decomposition.

3.1 Connection with LDS Decomposition

Let $X \subseteq V$ be a non-empty subset of the vertex set. Define the *induced edges* of X as $E(X) = \{(x, y) \in E \mid x, y \in X\}$, and the *induced subgraph* of X in G is $(X, E(X))$. The *density* of X is characterized by $\rho(X) = |E(X)|/|X|$. The subset $X \subseteq V$ that maximizes the density $\rho(X)$ is recognized as the *densest subgraph* of G .

For two vertex subsets X and Y with $X \cap Y = \emptyset$, we define the *cross edges* between X and Y as $E_{\times}(X, Y) = \{(x, y) \in E \mid x \in X, y \in Y\}$, and the *additional edges* from X with respect to Y as $E_{\Delta}(X, Y) = E(X) \cup E_{\times}(X, Y)$, representing the increment in edges upon incorporating X into Y . Based on this concept, we define the outer density and LDS [54] as follows.

DEFINITION 5. (Outer density) If $X \cap Y = \emptyset$, the outer density of a non-empty vertex set X with respect to set Y is defined as $\rho(X, Y) \triangleq |E_{\Delta}(X, Y)|/|X|$. If $X \cap Y \neq \emptyset$, the focus is on the subset of X not included in Y , i.e., $X \setminus Y$. In this case, we define $E_{\times}(X, Y) \triangleq E_{\times}(X \setminus Y, Y)$, $E_{\Delta}(X, Y) \triangleq E_{\Delta}(X \setminus Y, Y)$, and $\rho(X, Y) \triangleq \rho(X \setminus Y, Y)$.

DEFINITION 6. (Locally-densest subgraph, LDS) Given an undirected graph G and a vertex set $W \subseteq V$, W is an LDS if there is no $X \subseteq W$ and $Y, Y \cap W = \emptyset$, such that $\rho(X, W \setminus X) \leq \rho(Y, W)$.

To provide a more intuitive understanding, we rephrase the definition of LDS as follows: W is locally-densest if and only if $W = \emptyset$, or $W = V$, or $\min_{X \subseteq W} \rho(X, W \setminus X) > \max_{Y \cap W = \emptyset} \rho(Y, W)$. In other words, the minimum-density subgraph X within W still has a higher outer density compared to the maximum-density subgraph Y outside W . This demonstrates that W is “locally”-densest. Tatti [54] proved that all LDSes in a graph exhibit a hierarchical property, i.e., for any two LDSes X and Y , either $X \subseteq Y$ or $Y \subseteq X$. Therefore, LDSes can form a nested chain, which is the LDS decomposition.

DEFINITION 7. (LDS decomposition) [54] Given an undirected graph $G = (V, E)$, denote all of its non-empty LDSes by $\{B_0, \dots, B_k\}$, such that $B_k \subsetneq B_{k-1} \subsetneq B_{k-2} \subsetneq \dots \subsetneq B_0$. As \emptyset is also an LDS, the LDS decomposition is defined as a set of LDSes, represented as $\mathcal{B} = \{B_{k+1} = \emptyset, B_k, \dots, B_0\}$.

By the LDS decomposition, like IDN, we can obtain a density value for each vertex, referred to as the *fractional dense number*.

DEFINITION 8. (Fractional dense number, FDN) For a vertex $u \in B_i \setminus B_{i+1}$, the FDN of u is defined as $r_u = \rho(B_i, B_{i+1})$.

Based on these definitions, Tatti [54] further proved several useful properties of the LDS decomposition as the following lemma.

LEMMA 2. [54] For a graph G and its LDS decomposition \mathcal{B} , the following properties hold: (1) the smallest non-empty and largest LDSes, namely B_k and B_0 , are the densest subgraph and V , respectively; (2) B_i exhibits strictly increasing density as i increases, i.e., for two vertices $u \in B_j \setminus B_{j+1}$ and $v \in B_i \setminus B_{i+1}$ with $j > i$, $r_u = \rho(B_j, B_{j+1}) > \rho(B_i, B_{i+1}) = r_v$ holds.

Density decomposition and LDS decomposition are studied separately within two different research communities [11, 23, 54]. It is unclear whether there exists a direct connection between density decomposition and LDS decomposition. In this work, we fill this gap and establish, for the first time, a very close relationship between density decomposition and LDS decomposition, as shown in Theorem 2 and Theorem 3. The two theorems reveal that: (1) every layer of density decomposition is an LDS; (2) each IDN is the round-up value of the FDN.

LEMMA 3. Given a layer R_k , for any two non-empty sets $S \subseteq R_k$ and $T \subseteq V \setminus R_k$, we have $\rho(S, R_k \setminus S) > k - 1 \geq \rho(T, R_k)$.

PROOF. By Theorem 1, for any $S \subseteq R_k$ and any T disjoint from R_k , we have $E_{\Delta}(S, R_k \setminus S) > (k-1) \cdot |S|$ and $E_{\Delta}(T, R_k) \leq (k-1) \cdot |T|$, thus $\rho(S, R_k \setminus S) > k - 1 \geq \rho(T, R_k)$. \square

THEOREM 2. Given an undirected graph G and its density decomposition $\mathcal{R} = \{R_k\}$, for all $k = 0, 1, \dots, R_k$ is an LDS.

PROOF. By Lemma 3, the theorem can be directly proven. \square

THEOREM 3. Given a graph and a vertex v , $\bar{r}_v = \lceil r_v \rceil$ holds.

PROOF SKETCH. For an arbitrary vertex $v \in V$, suppose $v \in R_i \setminus R_{i+1}$ (i.e., $\bar{r}_v = i$) and $v \in B_j \setminus B_{j+1}$ (i.e., $r_v = \rho(B_j, B_{j+1})$). Let LDSes $B_{u+1} = R_{i+1}$ and $B_l = R_i$. In the proof of Lemma 3, considering respectively $i = k$, $S = B_l \setminus B_{l+1}$ and $i = k - 1$, $T = B_u \setminus B_{u+1}$, we get $i \geq \rho(B_u, B_{u+1}) \geq \rho(B_j, B_{j+1}) \geq \rho(B_l, B_{l+1}) > i - 1$. Thus $\bar{r}_v = i = \lceil \rho(B_j, B_{j+1}) \rceil = \lceil r_v \rceil$. \square

Note that the number of LDSes can be extensive and only be bounded by $O(n)$ [54]. In contrast, the density decomposition groups LDSes by a round-up relationship, with the number of non-trivial layers equaling the pseudoarboricity p , typically much smaller than n . This characteristic ensures that the density decomposition is not only computationally efficient but also retains a significant amount of information about the density structure of the graph.

Discussion: Existing algorithms for LDS decomposition. Based on Theorem 2 and Theorem 3, existing algorithms for computing LDS decomposition [23, 30, 54] can be applied to compute density decomposition by rounding up FDN to obtain IDN. Furthermore, building on these two theorems, we can also propose novel algorithms tailored for density decomposition. These novel algorithms are more efficient than the methods for computing LDS decomposition, as confirmed in our experiments.

3.2 Connection with Core Decomposition

Here, we reveal the relationship between the density decomposition and the core decomposition. The definitions of k -core and core decomposition are given as follows.

DEFINITION 9. (k -core and core decomposition) Given an undirected graph G and an integer k , the k -core of G , denoted by C_k , is the maximal subgraph in which every vertex has a degree of at least k . The core decomposition of G is a set of subgraphs, denoted by $C = \{C_k\}$, containing all k -cores of G .

As mentioned in Section 1, core decomposition may not always accurately reflect the density structure of a graph. However, we find that it can serve as a 2-approximation of density decomposition, according to the following theorems.

THEOREM 4. Given an undirected graph $G = (V, E)$, for $k = 0, 1, \dots, p$, $\rho(C_k, C_{k+1}) \geq \rho(R_k, R_{k+1})/2$.

PROOF. Let \vec{G} be an arbitrary egalitarian orientation of G . Denote $C_k \setminus C_{k+1}$ as C and $R_k \setminus R_{k+1}$ as R . We have $\rho(C_k, C_{k+1}) = \frac{|E_{\Delta}(C_k, C_{k+1})|}{|C|} \geq \frac{\sum_{u \in C} d_u(C_k)}{2|C|} \geq \frac{k}{2} \geq \frac{\sum_{u \in R} d_u(\vec{G})}{2|R|} = \frac{\rho(R_k, R_{k+1})}{2}$. \square

THEOREM 5. (Sandwich Theorem) Given an undirected graph G , for any non-negative integer k , we have $C_{2k} \subseteq R_k \subseteq C_k \subseteq R_{\lceil k/2 \rceil}$.

PROOF SKETCH. Let \vec{G} be an egalitarian orientation of G . To prove $R_k \subseteq C_k$, for any vertex $u \in R_k$, we either have $d_u(\vec{G}) \geq k$, or $d_u(\vec{G}) = k - 1$ and having an out-edge. So, in R_k , every vertex has a degree of at least k , and thus in C_k . To prove $C_k \subseteq R_{\lceil k/2 \rceil}$, when $C_k \neq \emptyset$, let u be an arbitrary vertex in C_k , below we show that $u \in R_{\lceil k/2 \rceil}$. Let $S \triangleq \{v | v \in C_k \text{ and } u \text{ can reach } v \text{ in } \vec{G}\}$. We have $\sum_{v \in S} d_v(\vec{G}) \geq |E_{\times}(S, C_k \setminus S)| + |E(S)| \geq \sum_{v \in S} d_v(C_k)/2 \geq k|S|/2$. Thus, at least one vertex $t \in S$ has indegree of at least $\lceil k/2 \rceil$ in \vec{G} . By definition, u can reach t and thus $u \in R_{\lceil k/2 \rceil}$.

Since $C_k \subseteq R_{\lceil k/2 \rceil} \Rightarrow C_{2k} \subseteq R_k, C_{2k} \subseteq R_k$ also holds. \square

THEOREM 6. *Given an undirected graph G , for any non-negative integer k such that $R_k \neq \emptyset, \rho(C_k) < \rho(R_k)$ holds, unless $C_k = R_k$.*

PROOF. When $C_k \neq R_k$, since $R_k \subseteq C_k$, we define $T = C_k \setminus R_k$. Then by Lemma 3, where let $S = R_k$, we have $\rho(C_k) = \rho(R_k \cup T) = \frac{|E(R_k)| + |E_{\Delta}(T, R_k)|}{|R_k| + |T|} < \frac{\rho(R_k) \cdot |R_k| + \rho(R_k) \cdot |T|}{|R_k| + |T|} = \rho(R_k)$. \square

Discussions. In contrast to core decomposition, density decomposition exhibits several advantages: (i) While k -cores are defined based on the degree, R_k fundamentally resembles LDS, as both are defined based on density, thus more accurately reflecting the density structure of graphs; (ii) Theoretically, as stated in Theorem 6, R_k is typically denser than C_k ; (iii) Practically, density decomposition proves more effective at identifying dense communities within graphs compared to core decomposition, as evidenced by multiple case studies in our experiments.

4 ALGORITHMS FOR STATIC GRAPHS

Before delving into static algorithms, we present Theorem 7, providing a method for obtaining a single layer of density decomposition.

THEOREM 7. *Given an orientation \vec{G} and a non-negative integer k , let $L = \{u | d_u(\vec{G}) < k - 1\}$ and $H = \{u | d_u(\vec{G}) > k - 1\}$. If there is no path starting from L and ending at H , then $R_k = H \cup \{u | u \text{ can reach a vertex in } H\}$.*

Based on Theorem 7, a layer R_k can be obtained by reversing all reversible paths from S to T , followed by a single execution of the BFS (Breadth-First Search) algorithm to identify $R_k = H \cup \{u | u \text{ can reach a vertex in } H\}$. The key to this method hinges on effectively reversing these paths from S to T . We first introduce an $O(m^2)$ -time complexity existing algorithm, Path [10, 11], which reverses these paths individually using multiple BFS operations. Then, to reduce the time complexity, we propose a novel and more efficient algorithm Flow, which utilizes the flow network technique to reverse these paths all at once. Compared to Path, Flow improves the time complexity from $O(m^2)$ to $O(m^{3/2} \cdot p)$. Building upon Flow, we develop the Flow++ algorithm, applying a divide-and-conquer strategy, further reducing the time complexity to $O(m^{3/2} \cdot \log p)$.

4.1 The Existing Path Algorithm

Borradaile [10, 11] proposed an algorithm for computing egalitarian orientation. Building upon this and our Theorem 7, we slightly modify it to develop an algorithm for computing density decomposition, denoted as Path, as shown in Algorithm 1. The Path algorithm sequentially computes each layer, starting from $R_0 = V$ and continuing until $R_k = \emptyset$. The method for computing each layer follows Theorem 7, where Path performs a single BFS to reverse

Algorithm 1: Path(\vec{G})

Input: An undirected graph G .
Output: Density decomposition $\mathcal{R} = \{R_k\}$.

- 1 Arbitrarily orient the edges in G to obtain \vec{G} ;
- 2 **foreach** $k = 0, 1, 2, \dots$ **do**
- 3 **while** *True* **do**
- 4 Try to find a reversible path $s \rightsquigarrow t$ using BFS algorithm, where
 $d_s(\vec{G}) < k - 1 < d_t(\vec{G})$;
- 5 **if** such path $s \rightsquigarrow t$ is found **then** reverse the path $s \rightsquigarrow t$;
- 6 **else** break the 'while' loop;
- 7 $H \leftarrow \{u \in \vec{G} | d_u(\vec{G}) > k - 1\}$;
- 8 $R_k \leftarrow H \cup \{u | u \text{ can reach a vertex in } H\}$;
- 9 **if** $R_k = \emptyset$ **then** break the 'foreach' loop;
- 10 **return** $\mathcal{R} = \{R_k\}$;

Algorithm 2: GetLayer(\vec{G}, k)

Input: An orientation $\vec{G} = (V, \vec{E})$ and an integer k .
Output: The updated \vec{G} and the layer R_k of the density decomposition.

- 1 $V' \leftarrow V \cup \{s, t\}$, where s is source and t is sink;
- 2 $d \leftarrow k - 1$;
- 3 **foreach** $\langle u, v \rangle \in \vec{E}$ **do**
- 4 Add arc $\langle u, v \rangle$ to A and let $c(u, v) \leftarrow 1$;
- 5 **foreach** $u, d_u(\vec{G}) < d$ **do**
- 6 Add arc $\langle s, u \rangle$ to A and let $c(s, u) \leftarrow d - d_u(\vec{G})$;
- 7 **foreach** $u, d_u(\vec{G}) > d$ **do**
- 8 Add arc $\langle u, t \rangle$ to A and let $c(u, t) \leftarrow d_u(\vec{G}) - d$;
- 9 Compute the maximum flow value f_{max} of (V', A, c) ;
- 10 **foreach** $\langle u, v \rangle \in \vec{E}$ **do** // Copy the residual network to \vec{G}
- 11 **if** $\langle u, v \rangle \in A$ is saturated **then** reverse the edge $\langle u, v \rangle \in \vec{E}$;
- 12 $R_k \leftarrow \{u \in V | d_u(\vec{G}) \geq d \text{ or } u \text{ can reach a vertex } v \text{ with } d_v(\vec{G}) \geq d\}$;
- 13 **return** (\vec{G}, R_k) ;

one reversible path at a time, with each BFS consuming $O(m)$ time (Lines 4-6). However, there may be a large number of reversible paths, making the time complexity of Path be as high as $O(m^2)$ [10]. Clearly, the high computational demand of the Path algorithm poses a significant challenge when handling large graphs.

4.2 A Re-orientation Network Flow Algorithm

As discussed above, the inefficiency of Path stems from reversing reversible paths one by one. A more effective approach involves employing the network flow technique to reverse multiple paths simultaneously. Therefore, we develop a novel algorithm based on the re-orientation network [6]. This method enables the simultaneous reversal of all reversible paths from L to H as specified in Theorem 7, significantly reducing the time cost.

Note that the re-orientation network flow technique was originally developed to minimize the indegree in an orientation [6], which differs essentially from density decomposition. To the best of our knowledge, we are the first to use the re-orientation network to compute density decomposition. Below, we introduce the re-orientation network.

DEFINITION 10. (Re-orientation network) [6] *Given an orientation $\vec{G} = (V, \vec{E})$ and an integer d , the re-orientation network is a weighted network with an additional source vertex s and sink vertex t where the weight of each edge denotes the capacity of the edge. Specifically, the re-orientation network with parameter d is defined as $(V \cup \{s, t\}, A, c)$ where (1) $\langle u, v \rangle \in A, c(u, v) = 1$, if*

Algorithm 3: Flow(\vec{G})

Input: An undirected graph G .
Output: Density decomposition $\mathcal{R} = \{R_k\}$.

- 1 Invoke a 2-approximation algorithm to obtain an approximate orientation \vec{G} ;
- 2 **foreach** $k = 0, 1, 2, \dots$ **do**
- 3 $(\vec{G}, R_k) \leftarrow \text{GetLayer}(\vec{G}, k)$;
- 4 **if** $R_k = \emptyset$ **then break**;
- 5 $\vec{G} \leftarrow$ the induced subgraph of R_k in \vec{G} ; // pruning strategy
- 6 **return** $\mathcal{R} = \{R_k\}$;

$\langle u, v \rangle \in E$; (2) $\langle s, u \rangle \in A, c(s, u) = d - d_u(\vec{G})$, if $d_u(\vec{G}) < d$; and
(3) $\langle u, t \rangle \in A, c(u, t) = d_u(\vec{G}) - d$, if $d_u(\vec{G}) > d$.

By Definition 10, the re-orientation network uses a parameter d to separate vertices based on their indegree. To compute a layer R_k , we set $d = k - 1$. Consequently, the source s connects to vertices with an indegree less than $k - 1$ (i.e., the set L in Theorem 7), while the sink t links to vertices with an indegree greater than $k - 1$ (i.e., the set H in Theorem 7). Upon completion of the maximum flow algorithm, no augmentation paths remain in the residual network, indicating no paths from s to t , namely L to H . Therefore, all reversible paths in Theorem 7 are reversed and the layer R_k can be easily obtained. Based on the above rationale, we devise the re-orientation network flow algorithm `GetLayer`, as shown in Algorithm 2, which admits an integer k and outputs R_k . Below, we give its correctness proof.

THEOREM 8. *Algorithm `GetLayer` correctly outputs R_k .*

PROOF SKETCH. By the above discussion, there is no path $s \rightsquigarrow t$ such that $d_s(\vec{G}) < k - 1$ and $d_t(\vec{G}) > k - 1$. As a result, by Lemma 7, the output R_k is correct. \square

The time complexity of Algorithm 2 depends on computing the maximum flow in the re-orientation network. It was shown that the re-orientation network maximum flow algorithm can be implemented with time complexity of $O(m^{3/2})$ and space complexity of $O(m)$ [7], and thus Algorithm 2 has the same complexity.

4.3 The Proposed Flow Algorithm

As shown in Theorem 8, each invocation of the `GetLayer` algorithm yields a single layer R_k . Therefore, we can substitute Lines 3-8 of `Path` with `GetLayer` to compute each R_k . Based on this, we propose the `Flow` algorithm, as depicted in Algorithm 3.

Algorithm `Flow` first invokes an existing 2-approximation algorithm to obtain a 2-approximation orientation in linear-time [25] (Line 1). Note that with such a 2-approximation initial orientation, the maximum flow computation can be faster than that with arbitrary initial orientation, because this approach is expected to reduce the number of augmentation paths in the maximum flow computation as indicated in [7]. Then, the `Flow` algorithm iteratively performs `GetLayer` from the bottom layer R_0 to yield each layer one by one (Lines 2-5). Furthermore, a pruning strategy is that, once R_k is computed, the algorithm can directly assign \vec{G} as the induced subgraph of R_k in \vec{G} by discarding the vertices in $V \setminus R_k$ (Line 5). Since the indegree of vertices in $V \setminus R_k$ is no larger than $k - 1$, they cannot reach vertices in R_k . In the following ‘foreach’ loop, the vertices in $V \setminus R_k$ are all connected to the source s , but the flow cannot enter R_k from $V \setminus R_k$. Thus, `Flow` only needs to compute R_{k+1} within R_k in the subsequent computation, which can significantly improve the efficiency. When k equals $p + 1$, R_k

Algorithm 4: Flow++(G)

Input: An undirected graph G .
Output: The egalitarian orientation \vec{G} and density decomposition $\{R_k\}$.

- 1 Invoke a 2-approximation algorithm to obtain an approximate orientation \vec{G} and a 2-approximation of pseudoarboricity \bar{p} ;
- 2 $R_{\bar{p}+1} \leftarrow \emptyset; R_0 \leftarrow V$;
- 3 `Divide`($R_{\bar{p}+1}, R_0$);
- 4 **return** ($\vec{G}, \{R_k\}$);

Function `Divide`(R_u, R_l)

- 6 **if** $u - l \leq 1$ or $R_u = R_l$ **then return**;
- 7 $k_u \leftarrow u, k_l \leftarrow l$;
- 8 **while** $k_u > k_l$ **do**
- 9 $k \leftarrow \lfloor (k_u + k_l + 1)/2 \rfloor$;
- 10 $(\vec{G}, R_k) \leftarrow \text{GetLayer++}(\vec{G}, k, R_u, R_l)$;
- 11 **if** $|E_\Delta(R_l, R_k)| < |E_\Delta(R_l, R_u)|/2$ **then** $k_l \leftarrow k$;
- 12 **else** $k_u \leftarrow k - 1$;
- 13 $k \leftarrow k_l$;
- 14 `Divide`(R_k, R_l);
- 15 $(\vec{G}, R_{k+1}) \leftarrow \text{GetLayer++}(\vec{G}, k + 1, R_u, R_l)$;
- 16 `Divide`(R_u, R_{k+1});

becomes an empty set, indicating that all layers of the density decomposition are obtained, thus the `Flow` algorithm terminates. The correctness of `Flow` can be guaranteed by Theorem 8.

For the complexity of `Flow`, since `Flow` invokes `GetLayer` $p + 2$ times, it can be easily derived that the time complexity and space complexity of `Flow` are respectively $O(p \cdot m^{3/2})$ and $O(m)$. Note that since $p \leq \sqrt{m}$ [17], the time complexity $O(p \cdot m^{3/2})$ is lower than the $O(m^2)$ time complexity of `Path`. Furthermore, p is often small in real-world graphs [7], thus the practical performance of `Flow` is typically much better than its worst-case time complexity.

4.4 The Proposed Flow++ Algorithm

In this subsection, we propose an improved algorithm `Flow++`. By using a novel divide-and-conquer technique, `Flow++` improves the time complexity of `Flow` from $O(p \cdot m^{3/2})$ to $O(\log p \cdot m^{3/2})$. The key idea of `Flow++` is that it partitions the graph into almost equal parts at each recursion and performs maximum flow computation recursively in each part. Recall that in `Flow`, we compute R_k within R_{k-1} instead of the whole graph G due to the hierarchical structure of density decomposition. Similar to this strategy, an improved strategy is to compute R_k between R_u and R_l , where $u > k > l$, which can intuitively further reduce the data scale. To construct such R_u and R_l , a divide-and-conquer manner can be employed to recursively compute all layers of density decomposition. Specifically, the recursive function, denoted as `Divide`(R_u, R_l), is designed to find all R_k for $u > k > l$. It first selects a parameter k where $u > k > l$ and divides the graph into two parts: $R_l \setminus R_k$ and $R_{k+1} \setminus R_u$. Then, `Divide`(R_k, R_l) and `Divide`(R_u, R_{k+1}) are separately invoked for deeper recursions.

Equipped with such a divide-and-conquer approach, the improved algorithm `Flow++` is proposed to compute the density decomposition, which is outlined in Algorithm 4. Specifically, Algorithm 4 starts by invoking an approximate algorithm to calculate a 2-approximate pseudoarboricity $\bar{p} \geq p$ (Line 1). Then, we can obtain two trivial layers: $R_{\bar{p}+1} = \emptyset$ and $R_0 = V$ (Line 2). The recursion begins with these two initial layers by calling `Divide`($R_{\bar{p}+1}, R_0$) (Line 3). During this invocation, a parameter k is selected to divide the graph (Lines 7-13). We use a binary search (Lines 7-13) to find the maximum k satisfying $|E_\Delta(R_l, R_k)| < |E_\Delta(R_l, R_u)|/2$. Thus,

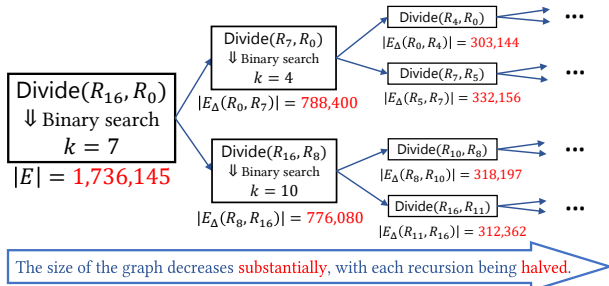


Figure 2: Illustration of the Flow++ algorithm on the Citeseer dataset ($|V| = 384,054$, $|E| = 1,736,145$).

for $k + 1$, we have $|E_{\Delta}(R_l, R_{k+1})| \geq |E_{\Delta}(R_l, R_u)|/2$, suggesting that $|E_{\Delta}(R_{k+1}, R_u)| = |E_{\Delta}(R_l, R_u)| - |E_{\Delta}(R_l, R_{k+1})| < |E_{\Delta}(R_l, R_u)|/2$. Therefore, both the edges $E_{\Delta}(R_l, R_k)$ and the edges $E_{\Delta}(R_{k+1}, R_u)$ are less than half of the edges $E_{\Delta}(R_l, R_u)$. This ensures the number of edges being divided by 2 in each depth of recursion. When $u - l \leq 1$ or $R_u = R_l$, there are no further subdivided layers between R_u and R_l , thus Divide can return directly. The correctness of the Flow++ algorithm can be easily derived by induction because in each recursion we can identify a R_k and no layer will be missed by the divide-and-conquer procedure. Thanks to our divide-and-conquer technique, the graph size is halved with each recursion. As a result, the computation costs exponentially decrease for the higher-level layers. This is why our Flow++ algorithm can achieve a lower time complexity compared to the Flow algorithm. Below, we use an example to illustrate how Flow++ works.

EXAMPLE 2. Assuming that the Flow++ algorithm processes the Citeseer dataset ($|V| = 384,054$, $|E| = 1,736,145$), it initially computes the 2-approximation pseudoarboricity $\bar{p} = 16$ and then invokes the Divide procedure for recursive computation. As shown in Fig. 2, in the first recursion round, i.e., $\text{Divide}(R_{16}, R_0)$, the algorithm employs binary search to determine $k = 7$ and subsequently divides the graph $R_0 \setminus R_{16}$ into two parts: $R_8 \setminus R_{16}$ and $R_0 \setminus R_7$. Both $|E_{\Delta}(R_0, R_7)|$ and $|E_{\Delta}(R_8, R_{16})|$ are less than $\frac{|E|}{2} = \frac{|E_{\Delta}(R_0, R_{16})|}{2}$, resulting in the partition of the graph into two smaller subgraphs. Subsequently, deeper recursion is independently conducted on the two smaller subgraphs, significantly reducing the graph scale $|E_{\Delta}(R_l, R_u)|$ which is halved in each recursion round. Therefore, the scale of the graph remains comparable to the original only at shallow recursion depths. As the recursion depth increases, the scale of the graph diminishes exponentially, thereby significantly improving efficiency.

The remaining challenge in the above divide-and-conquer algorithm lies in implementing the computation within $R_l \setminus R_u$. We find that we can slightly modify the re-orientation network flow algorithm to achieve this goal, obtaining the improved version GetLayer++ of GetLayer. The GetLayer++ algorithm inputs an integer k and two layers R_u, R_l , satisfying $u \geq k \geq l$. It outputs the updated \vec{G} and R_k with a time complexity of $O(|E_{\Delta}(R_l, R_u)|^{3/2})$. To implement it, we need to consider not only the subgraph induced by $R_l \setminus R_u$ in \vec{G} but also the cross edges $E_{\times}(R_l, R_u)$. In GetLayer++, arcs from the source and sink to vertices in only $R_l \setminus R_u$ are connected (the other vertices are pruned). However, for computing the indegree of the vertices in $R_l \setminus R_u$, the edges in $E_{\times}(R_l, R_u)$ are also considered. The reasons why GetLayer++ can ensure the correctness are as follows. Because R_l and R_u are either previously

computed by the re-orientation network or are the trivial subgraphs \emptyset or V . According to the properties of the re-orientation network, the edges in $E_{\times}(V \setminus R_u, R_u)$ all point towards $V \setminus R_u$, and the edges in $E_{\times}(V \setminus R_l, R_l)$ all point towards $V \setminus R_l$. This indicates that at this point when the GetLayer++ algorithm is called again for further computation, there will be no flow from $V \setminus R_u$ to R_u , and no flow from $V \setminus R_l$ to R_l . As a result, we can safely ignore $V \setminus R_l$ and R_u , which do not have any flow passing through them, and only connect the vertices in $R_l \setminus R_u$ to the source and sink.

Clearly, with this divide-and-conquer implementation, the size of the network can be bounded by $E_{\Delta}(R_l, R_u)$ in each recursion. Thus, by [7], the time complexity of $\text{GetLayer++}(\vec{G}, k, R_u, R_l)$ is $O(|E_{\Delta}(R_l, R_u)|^{3/2})$. Based on this, we can derive the total time and space complexity of Flow++ as follows.

THEOREM 9. The time complexity of Algorithm 4 is $O(m^{3/2} \cdot \log p)$, and the space complexity is $O(p \cdot n)$.

PROOF SKETCH. The time complexity of $\text{Divide}(R_l, R_u)$ consists of two parts: the complexity of binary search and the complexity of the deeper recursion $\text{Divide}(R_k, R_l)$ and $\text{Divide}(R_u, R_{k+1})$. For binary search, it invokes GetLayer++ $O(\log p)$ times, so the time complexity of the binary search in $\text{Divide}(R_u, R_l)$ is $O(|E_{\Delta}(R_l, R_u)|^{3/2} \cdot \log p)$. For the deeper recursion, the selection of k guarantees that both $|E_{\Delta}(R_l, R_k)|$ and $|E_{\Delta}(R_{k+1}, R_u)|$ are less than $|E_{\Delta}(R_l, R_u)|/2$, ensuring that the data scale is divided by 2 in each recursive layer. According to the master theorem for divide-and-conquer algorithms [4], the time complexity of deeper recursion in $\text{Divide}(R_u, R_l)$ is less than that of the binary search. Hence, the time complexity of Flow++ is the time complexity of the binary search in $\text{Divide}(R_{\bar{p}+1}, R_0)$, which is $O(m^{3/2} \cdot \log p)$. For the space complexity, storing the space of R_u and R_l for each recursion depth takes $O(p \cdot n) \supseteq O(m)$ space [42]. So the space complexity of Flow++ is $O(p \cdot n)$. \square

Note that since p is small in real-world graphs, $\log p$ is a very small constant. For example, for most real-world graphs where $p \leq 1024$, we have $\log_2 p \leq 10$. Consequently, the complexity of our algorithm is near to $O(m^{3/2})$, making it highly efficient for processing real-world graphs.

5 DYNAMIC MAINTENANCE ALGORITHMS

When the graph is updated by an edge insertion or deletion, a straightforward algorithm to maintain the density decomposition is to invoke the Flow++ algorithm to re-compute the density decomposition from scratch. Clearly, such an approach is costly. To improve the efficiency, in this section, we develop several novel algorithms to maintain the density decomposition when the graph is updated by an edge insertion or deletion. To avoid confusion, we assume that the notations, including IDN \bar{r} , indegree d , and R_k , all denote their values *before* edge updated unless otherwise specified.

5.1 Density Decomposition Update Theorem

In this subsection, we establish a density decomposition update theorem, based on which we can devise the maintenance algorithms.

THEOREM 10. (Density decomposition update theorem) When inserting (resp. deleting) an edge (u, v) in the undirected graph G , assuming $\bar{r}_v \leq \bar{r}_u$, only the IDNs of vertices whose IDN equals \bar{r}_v , i.e., vertices in $R_{\bar{r}_v} \setminus R_{\bar{r}_v+1}$, may change, and it can only increase by 1 (resp. decrease by 1).

Algorithm 5: Insert($\vec{G}, \vec{r}, (u, v)$)

Input: The egalitarian orientation \vec{G} , the IDNs of all vertices \vec{r} , and the edge (u, v) to be inserted.

Output: The updated egalitarian orientation \vec{G} and IDNs \vec{r} .

```
1 Suppose  $\vec{r}_v \leq \vec{r}_u$ , otherwise swap the input edge  $(u, v)$ ;  
2 if  $d_v = \vec{r}_v - 1$  then  $\vec{G} \leftarrow \vec{G} \cup \langle u, v \rangle$ ;  
3 else //  $d_v = \vec{r}_v$   
4    $\vec{G} \leftarrow \vec{G} \cup \langle u, v \rangle$ ;  
5   if there is a reversible path  $s \rightsquigarrow v$ , where  $d_s = \vec{r}_v - 1$  then  
6     Reverse the path;  
7   else  
8     foreach  $w \in \{w | \vec{r}_w = \vec{r}_v \text{ and } w \text{ can reach } v\}$  do  $\vec{r}_w \leftarrow \vec{r}_w + 1$ ;  
9      $\vec{r}_v \leftarrow \vec{r}_v + 1$ ;  
10 return  $(\vec{G}, \vec{r})$ ;
```

Theorem 10 indicates that maintaining density decomposition is highly localized. That is, we only need to consider the vertices within $R_{\vec{r}_v} \setminus R_{\vec{r}_v+1}$. Besides, given that the density decomposition is defined based on an egalitarian orientation, the key to updating the density decomposition lies in maintaining the egalitarian orientation, building upon which we develop several efficient density decomposition maintenance algorithms in the following.

5.2 The Proposed Insertion Algorithm

Based on Theorem 10, we propose an insertion algorithm, called Insert, to handle the case of edge insertion. The detailed implementation of our algorithm is shown in Algorithm 5. The main idea of Insert is as follows. By Theorem 10, Insert needs to consider three different cases: (1) $d_v = \vec{r}_v - 1$ (Line 2); (2) $d_v = \vec{r}_v$ and there exists $s \rightsquigarrow v$ (Lines 3-6); (3) $d_v = \vec{r}_v$ and there is no $s \rightsquigarrow v$ (Lines 3-4 and Lines 7-9). In case (1), the Insert algorithm directly inserts $\langle u, v \rangle$. Following this operation, \vec{G} is still an egalitarian orientation, and the IDNs do not need to be updated. In case (2), Insert also inserts $\langle u, v \rangle$, but a reversible path $s \rightsquigarrow v$ is created, making \vec{G} non-egalitarian. Consequently, Insert reverses this path, restoring \vec{G} to an egalitarian state. At this point, the IDNs also do not change. Note that for case (2), we can perform once BFS from the vertex v , traversing in the opposite direction of the edges, to find the reversible path $s \rightsquigarrow v$ (Lines 5-6). In case (3), where no reversible path is found, \vec{G} remains an egalitarian orientation, but the IDNs of some vertices need to be increased by 1 according to Theorem 10. Similarly, in this case, invoking once BFS from v can identify the set $\{w | \vec{r}_w = \vec{r}_v \text{ and } w \text{ can reach } v\}$ (Line 8). Below, we prove that Insert correctly maintains the density decomposition (i.e., maintains the IDNs), and analyze its complexity.

THEOREM 11. *Algorithm 5 Insert correctly maintains an egalitarian orientation and all IDNs for all vertices.*

THEOREM 12. *The worst-case time complexity of Algorithm 5 is $O(\sum_{u \in R_{\vec{r}_v} \setminus R_{\vec{r}_v+1}} d_u(G))$.*

As shown in Theorem 12, the time complexity of Algorithm 5 is clearly bounded by $O(m)$. Furthermore, its complexity only depends on a typically small subgraph $R_{\vec{r}_v} \setminus R_{\vec{r}_v+1}$, indicating that the algorithm is highly local.

5.3 The Proposed Deletion Algorithm

In this subsection, we proposed the Delete algorithm to maintain the density decomposition for handling edge deletion, as outlined

Algorithm 6: Delete($\vec{G}, \vec{r}, (u, v)$)

Input: The egalitarian orientation \vec{G} , the IDNs of all vertices \vec{r} , and the edge (u, v) need to be deleted.

Output: The updated egalitarian orientation \vec{G} and IDNs \vec{r} .

```
1 Suppose  $(u, v)$  is oriented as  $\langle u, v \rangle$ , otherwise swap the input edge  $(u, v)$ ;  
2 if  $d_v = \vec{r}_v - 1$  then  
3   There must be a reversible path  $v \rightsquigarrow t$ , where  $d_t = \vec{r}_v$ ;  
4   Reverse the path;  
5  $\vec{G} \leftarrow \vec{G} - \langle u, v \rangle$ ;  
6 Let  $\vec{r}_0 \leftarrow \vec{r}_v$ ;  
7 Let  $S \leftarrow \{w | \vec{r}_w = \vec{r}_0, \text{ and } d_w = \vec{r}_0 \text{ or } w \text{ can reach an } \vec{r}_0\text{-indegree vertex}\}$ ;  
8 forall  $w$  where  $\vec{r}_w = \vec{r}_0$  and  $w \notin S$  do  
9    $\vec{r}_w \leftarrow \vec{r}_w - 1$ ;  
10 return  $(\vec{G}, \vec{r})$ ;
```

in Algorithm 6. Suppose that the deleted edge (u, v) is oriented as $\langle u, v \rangle$ (Line 1). According to Lemma 1, the inequality $\vec{r}_v \leq \vec{r}_u$ holds, similar to the Insert algorithm. The Delete algorithm addresses two cases for maintaining the egalitarian orientation: (1) $d_v = \vec{r}_v$ (not satisfying the ‘if’ condition of Line 2); (2) $d_v = \vec{r}_v - 1$ (satisfying the ‘if’ condition). In case (1), Delete directly deletes $\langle u, v \rangle$ (Line 5), and \vec{G} is still an egalitarian orientation. In case (2), by the definition of density decomposition, a reversible path $v \rightsquigarrow t$ must exist (Line 3). This ensures that Delete does not encounter the scenario of being unable to find a reversible path, as opposed to Insert. After reversing $v \rightsquigarrow t$, case (2) transforms into case (1). Then, $\langle u, v \rangle$ can be deleted and \vec{G} is restored to the egalitarian orientation.

With the egalitarian orientation, the Delete algorithm starts maintaining the IDNs of vertices. Unfortunately, unlike Insert, which can update IDNs directly, there is no direct method to update IDNs in the deletion case. Therefore, Delete needs to perform Lines 6-9 to update IDNs. Since \vec{r}_v may change, to avoid ambiguity, we use \vec{r}_0 to denote the value of \vec{r}_v before updating (Line 6). By the definition of density decomposition, Delete employs a single BFS algorithm to determine the set of vertices S that do not require IDNs updates (Line 7). It then decreases the IDNs of vertices that were originally equal to \vec{r}_0 but are not in the set S (Lines 8-9). Finally, the Delete algorithm terminates, outputting the egalitarian orientation and the IDNs of all vertices. The following theorem shows the correctness of Algorithm 6.

THEOREM 13. *Algorithm 6 correctly updates the egalitarian orientation and the IDNs of all vertices.*

THEOREM 14. *The worst-case time complexity of Algorithm 6 is $O(\sum_{u \in R_{\vec{r}_v} \setminus R_{\vec{r}_v+1}} d_u(G))$.*

Similar to the Insert algorithm, the Delete algorithm is also local, and the worst-case time complexity is clearly bounded by $O(m)$.

5.4 An Improved Deletion Algorithm

Despite exhibiting the same time complexity, in our experiments, the Delete algorithm runs around three orders of magnitude slower than Insert. The main issue is, to update the IDNs of vertices, Insert performs BFS only within $\{w | \vec{r}_w = \vec{r}_v \text{ and } w \text{ can reach } v\} \subseteq R_{\vec{r}_v} \setminus R_{\vec{r}_v+1}$, whereas Delete necessitates considering the entire $R_{\vec{r}_v} \setminus R_{\vec{r}_v+1}$. In this subsection, we devise an improved algorithm, called Delete++ (Algorithm 7), to further enhance the efficiency.

The rationale and intuition behind Delete++ are as follows. If the IDN of a vertex w decreases due to the deletion of $\langle u, v \rangle$, by Definition 3, w must have been capable of reaching v prior to the

Algorithm 7: Delete++($\vec{G}, \bar{r}, (u, v)$)

Input: The egalitarian orientation \vec{G} , the IDNs of all vertices \bar{r} , and the edge (u, v) to be deleted.

Output: The updated egalitarian orientation \vec{G} and IDNs \bar{r} .

- 1 Perform Lines 1-4 of Delete;
- 2 Let $P \leftarrow \{w | w \text{ can reach } v\} \cap (R_{\bar{r}_v} \setminus R_{\bar{r}_{v+1}})$;
- 3 $\vec{G} \leftarrow \vec{G} - (u, v)$;
- 4 **foreach** $s \in P$ **and** $d_s = \bar{r}_v - 1$ **do**
- 5 **if** the SCC containing s has not been computed **then**
- 6 Let $S \leftarrow \{s\} \cup \{w | s \text{ can reach } w\} \cap (R_{\bar{r}_v} \setminus R_{\bar{r}_{v+1}})$;
- 7 Compute all SCCs in the set S , while recording whether they can reach a \bar{r}_v -indegree vertex;
- 8 **if** the SCC containing s cannot reach a vertex with \bar{r}_v indegree **then**
- 9 $\bar{r}_s \leftarrow \bar{r}_s - 1$;
- 10 **return** (\vec{G}, \bar{r});

deletion. Otherwise, w would have possessed the ability to reach another vertex with an \bar{r}_v -indegree distinct from v before the deletion, thereby preventing its IDN from decreasing. Therefore, we can calculate the set $P \subseteq R_{\bar{r}_v} \setminus R_{\bar{r}_{v+1}}$, which includes the vertices capable of reaching v (Line 2). The vertices where the IDNs may be reduced are specifically within the set P , allowing us to narrow down the scope for updating the IDNs to P rather than $R_{\bar{r}_v} \setminus R_{\bar{r}_{v+1}}$.

Subsequently, we turn our focus to updating the IDNs of the vertices in P . Recall that the Delete algorithm identifies all vertices capable of reaching the vertices with \bar{r}_v -indegree, and then reduces the IDNs of vertices rather than these vertices. Such an approach, however, cannot be directly applied to update the IDNs within P . This is because, if a vertex in P cannot reach \bar{r}_v -indegree vertices within P , it does not necessarily imply that its IDN needs to be decreased since it might still have the possibility of reaching \bar{r}_v -indegree vertices outside of P .

A basic idea is performing a BFS individually on each vertex in P to determine its reachability to vertices with an indegree of \bar{r}_v . But it is impractical due to massive redundant re-computations. The key to solving this problem lies in reducing the number of BFS operations instead of “individually” searching for each vertex. To achieve this, we leverage a classic concept of Strongly Connected Component [53] (SCC). Once whether a vertex s can reach \bar{r}_v -indegree vertices is confirmed, vertices within the same SCC as s also share the same reachability, eliminating the need to invoke the BFS algorithm on these vertices again.

In the following, we show the correctness of the Delete++ algorithm. Since Line 1 and Line 3 of Delete++ is the same as the method for maintaining egalitarian orientation in Delete, Delete++ can correctly update egalitarian orientation. Next, we demonstrate that Delete++ can correctly update the IDNs of vertices.

THEOREM 15. *Algorithm 7 correctly updates the IDNs of vertices.*

We employ the linear-time Tarjan algorithm [53] to compute SCCs, and thus the complexity of Delete++ is equivalent to that of Delete. However, in contrast to Delete, which accounts for the entire set $R_{\bar{r}_v} \setminus R_{\bar{r}_{v+1}}$, Delete++ focuses solely on the neighborhood P of vertex v . Consequently, the actual running time of Delete++ proves notably faster than that of Delete in our experiments.

6 EXPERIMENTS

Different algorithms. For static graphs, we compare our proposed algorithms Flow (Algorithm 3) and Flow++ (Algorithm 4)

Table 2: Datasets statistics

Name	Type	n	m
DBLP	co-authorship network	317,081	1,049,866
Citeseer	citation network	384,414	1,736,145
Yahoo	lexical network	653,261	2,931,698
Skitter	internet	1,694,617	11,094,209
Weibo	social network	58,655,850	261,321,033
UKlink	web graph	18,483,187	261,787,258
Twitter	social network	20,826,113	294,585,816
Wiki	web graph	13,593,033	334,591,525

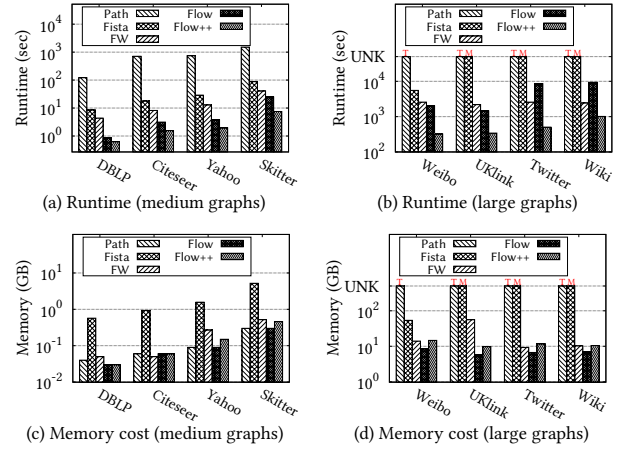


Figure 3: Results of density decomposition on static graphs.

with three baselines: Path (Algorithm 1) [11], FW [23], and Fista [30]. Here, Flow, Flow++, and Path are implemented by us. FW [23] and Fista [30] are the state-of-the-art algorithms for LDS decomposition. By our results established in Section 3, these algorithms can also be used to compute the density decomposition, thus we also include them as baselines. For dynamic graphs, we implement three proposed algorithms Insert (Algorithm 5), Delete (Algorithm 6), and Delete++ (Algorithm 7). Since there is no algorithm that can maintain density decomposition, we use the Flow++ algorithm for re-computing the density decomposition as a baseline.

All algorithms are implemented using C++ and compiled with the GCC compiler, employing the O3 optimization. All experiments are conducted on a PC operating a Linux system, equipped with a 2.2GHz AMD 3990X 64-Core CPU and 128GB memory.

Datasets. We use 4 medium graphs and 4 large graphs in our experiments which are downloaded from the Network Repository [44] and the Koblenz Network Collection (<http://konect.cc/>). The detailed information is provided in Table 2.

6.1 Performance Studies

Exp-1: Runtime of various density decomposition algorithms.

Fig. 3a and Fig. 3b depict the results of different algorithms for computing the density decomposition. In Fig. 3b, there are two cases designated as ‘UNK’ (unknown) for runtime: (1) the algorithm’s runtime exceeds our time constraint of 50,000 seconds, labeled as ‘T’; (2) the algorithm’s memory usage exceeds 128GB, labeled as ‘M’. From Fig. 3a and Fig. 3b, Path unsurprisingly exhibits the longest execution time, and even fails to complete the computation within the time constraint on all large graphs. In contrast, our proposed algorithms, Flow and Flow++, not only successfully compute the density decomposition for all datasets but are also at least 1

Table 3: Number of layers of different decompositions

“DD” denotes density decomposition.
“LDS” represents locally-densest subgraph decomposition.

Datasets	DD	LDS	Datasets	DD	LDS
DBLP	59	1,088	Weibo	168	5,609
Citeseer	16	1,435	UKlink	474	40,875
Yahoo	26	1,376	Twitter	840	11,949
Skitter	92	3,493	Wiki	602	23,203

and 2 orders of magnitude faster than Path respectively, exhibiting significant superiority. In addition, Flow++ shows remarkable enhancements, particularly on large graphs, where efficiency is improved by about an order of magnitude compared to Flow. These results confirm our theoretical analysis in Section 3.

When comparing the two LDS decomposition algorithms, we find that FW outperforms Fista over all datasets. The reason could be that Fista was originally designed as an approximate LDS decomposition algorithm, relying on a peeling technique [30], and it may not be well-suited for exact LDS decomposition. Compared to the LDS decomposition algorithms, our best algorithm Flow++ is consistently faster on all datasets, and can still achieve one order of magnitude speedup on large datasets. For instance, on Weibo, the runtime of Flow++ and FW are 190.42 seconds and 2,575.96 seconds respectively, representing a 13x acceleration. This result further demonstrates the high efficiency of our Flow++ algorithm.

Exp-2: Memory overheads of different algorithms. The memory usages of different algorithms are shown in Fig. 3c and Fig. 3d. Both the Path and Flow algorithms have lower memory costs, which are consistent with the theoretical analysis of the space complexity as $O(m)$. The algorithms Flow++, FW, and Fista all involve a recursive procedure, which requires auxiliary space usage. Among them, Flow++ needs to store each layer R_k , and since the number of layers is small as shown in Table 3, the memory cost of Flow++ remains modest. On all datasets, Flow++ requires no more than 15GB of memory. In the case of FW and Fista, their memory costs depend to a large extent on the effectiveness of the approximation algorithms that they employ. The approximation algorithm in FW is more effective, thus it often consumes less memory than Fista. Nevertheless, on UKlink, FW algorithm uses significantly higher memory compared to Flow and Flow++, indicating lower robustness in terms of memory cost for the FW algorithm. These results confirm that the memory usage of our algorithms is acceptable with efficient computational performance.

Exp-3: Number of layers of different decompositions. The number of layers of density decomposition and LDS decomposition are presented in Table 3. For density decomposition, the number of layers is equal to the pseudoarboricity plus 2, which is typically small in real-world graphs [35]. However, the LDS decomposition yields a considerably larger number of layers, reaching 40,875 layers on UKlink. This multitude of layers often leads to high computation time. Moreover, such a fine partition of layers may inevitably lead to unnecessary separation of tightly connected subgraphs, thus reducing the performance of identifying dense subgraphs for real applications. These results confirm our analysis in Sections 1 and 3.

Exp-4: Runtime of the various maintenance algorithms. We randomly delete and insert 10,000 edges from each graph dataset to evaluate different maintenance algorithms. The total time required to process these 10,000 deleted or inserted edges is shown in Fig. 4. For the baseline algorithm, since both insertion and deletion require invoking Flow++ to recompute, the runtime for deletion and insertion is the same, thus we only plot one bar in Fig. 4. As can be

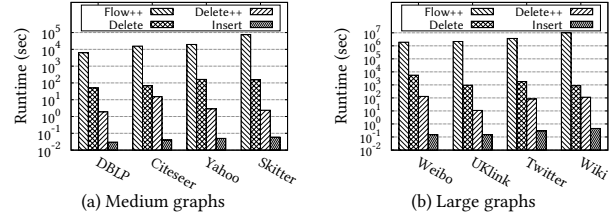


Figure 4: Runtime of various maintenance algorithms

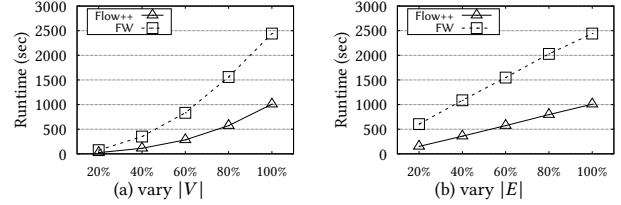


Figure 5: Scalability test of static algorithms on Wiki.

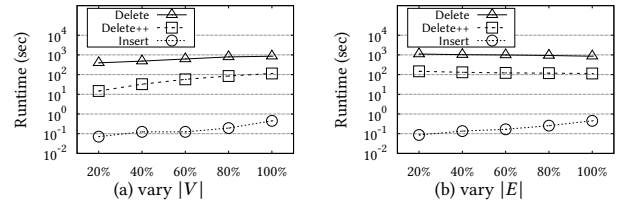


Figure 6: Scalability test of dynamic algorithms on Wiki.

seen, both Delete and Delete++ are around 3-5 orders of magnitude faster than the baseline algorithm (Flow++), respectively. Furthermore, Delete++ significantly outperforms Delete on all datasets as expected. For example, on UKlink, Delete consumes 943.62 seconds to maintain the density decomposition, while Delete++ only takes 10.98 seconds, achieving an 85x acceleration. These results confirm our analysis in Section 5. For edge insertion, Insert is substantially faster than Flow++, improving by about 5-7 orders of magnitude. When applied to large graphs such as Wiki with hundreds of millions of edges, handling the insertion of 10,000 edges only requires 0.4 seconds. These results further demonstrate the high efficiency of our algorithm. Additionally, it is worth noting that the runtime of our insertion algorithm is often orders of magnitude faster than that of the deletion algorithms. This significant efficiency stems from the fact that (as analyzed in Section 5) updating the IDNs during insertion is much easier compared to the deletion case.

Exp-5: Scalability test. In this experiment, we evaluate the scalability of static and dynamic algorithms. For static algorithms, we generate 8 subgraphs for each dataset by randomly sampling the vertex set V and edge set E , and then perform Flow++ and FW on these subgraphs. Results from our largest dataset, Wiki, are presented in Fig. 5, and similar patterns are observed across other datasets. As expected, the runtime of FW increases sharply with the data scale, in contrast to the more gradual increase observed with Flow++. Moreover, Flow++ outperforms FW at all data scales, consistent with previous observations. These findings demonstrate the high scalability of the Flow++ algorithm.

For dynamic algorithms, we compare the running time of Delete, Delete++, and Insert in randomly processing 10,000 edges, with the results presented in Fig. 6. In contrast to static algorithms, where

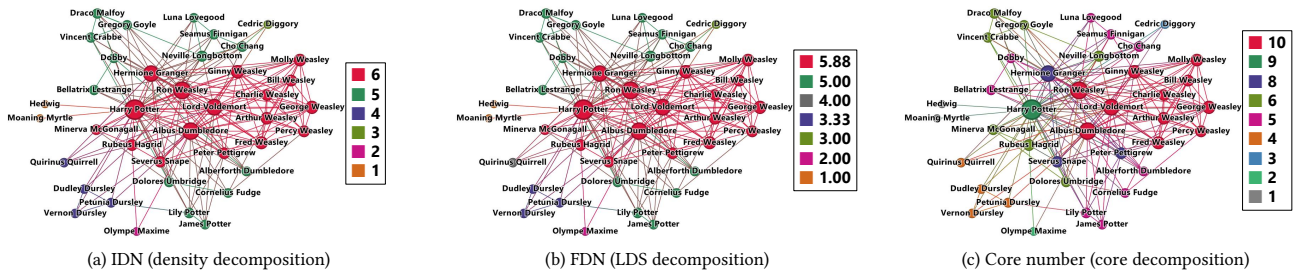


Figure 7: Dense subgraph decompositions of Harry Potter character relationship network HPC.

runtime significantly escalates with graph size, the runtime of dynamic algorithms exhibits notable insensitivity to graph size. This insensitivity arises because the complexity of our proposed dynamic algorithms is $O(\sum_{u \in R_{F_d} \setminus R_{F_{v+1}}} d_u(G))$ (details in Section 5), which is independent of graph size, theoretically ensuring the scalability of our algorithms. Moreover, Delete++ consistently outperforms Delete by approximately an order of magnitude across all graph sizes, and Insert maintains high efficiency with a running time of less than one second for all graph sizes. These results demonstrate the high scalability of our proposed maintenance algorithms.

6.2 Case Studies

Harry Potter characters relationship network. Here we conduct a case study on the Harry Potter characters network¹, denoted by HPC. We evaluate the effectiveness of three graph decompositions, including density decomposition, LDS decomposition, and core decomposition. The IDNs, FDNs, and core numbers of all characters in HPC are illustrated in Fig. 9, with vertex sizes corresponding to their degrees. From Fig. 7a and Fig. 7b, we can see that the density decomposition and LDS decomposition on HPC exhibit a high similarity, with the only distinction being that in the LDS decomposition, Quirinus Quirrell, one of the four characters with an IDN of 4, is separated from the other three characters. This result indicates that the density decomposition can well approximate the LDS decomposition. Moreover, in this case study, there seems no need to excessively differentiate the characters in fine detail. Whether a character with FDN of 4.00 or 3.33 can both be considered as similar important secondary characters. This result suggests that a moderately coarse-grained partition achieved by density decomposition often suffices, eliminating the need for an excessively fine-grained partition. Additionally, decomposing the graph into finer layers generally requires more computational time. On HPC, invoking Flow++ for density decomposition takes 0.048ms, while performing FW to compute LDS decomposition consumes 0.918ms. In summary, density decomposition can achieve similar results as the LDS decomposition, while significantly reducing computational time and the number of redundant dense subgraph layers.

When comparing the density decomposition (Fig. 7a) and the core decomposition (Fig. 7c), a noticeable disparity emerges. The most notable difference is that the top layer of density decomposition is divided into four layers in the core decomposition. The main character, Harry Potter, is not included in the top layer (i.e., 10-core), which also exhibits a clearly lower density in comparison to the other two density-aware decomposition methods. These results confirm our analysis in Sections 1 and 3.

Graph and Digraph Glossary network. Below, we conduct a case study on the Graph and Digraph Glossary network², denoted as GDG. This network summarizes a range of graph theory terms used in a course about Euler circuits and coloring problems. Each term within the network is depicted as a node, with edges illustrating the definition relationships among these terms. The density decomposition, LDS decomposition, and core decomposition on GDG are illustrated in Fig. 8. As shown in Fig. 8a, the density decomposition of GDG organizes the network into three strata corresponding to high, medium, and low densities. The top layer, R_3 , incorporates frequently used terms such as ‘Graph’ and ‘Vertex’, along with course-specific terms like ‘ k -colorable’, suggesting that R_3 comprises the fundamental terminology for the course. In Fig. 8b, the LDS decomposition, based on the density decomposition, further partitions GDG. This decomposition notably segments the low-density area from Fig. 8a into several sub-sections with FDN values of 1.50, 1.00, 0.67, and 0.50. However, this granular subdivision may be deemed superfluous, as the terms in these parts could be considered equally non-essential. Concerning core decomposition presented in Fig. 8c, it retains the tri-layer structure of the density decomposition but differentiates itself by incorporating some medium-density nodes from $R_2 \setminus R_3$ into the top layer C_3 . This inclusion results in some relatively non-essential terms like ‘Closure’ being unreasonably considered as essential terms. Moreover, the density of R_3 , $\rho(R_3) = 2.40$, surpasses that of C_3 , $\rho(C_3) = 2.21$. For the medium-density vertices added to C_3 , i.e., $C_3 \setminus R_3$, the outer density $\rho(C_3 \setminus R_3, R_3)$ is only 2, indicating that the core decomposition may inappropriately elevate the status of lower-density terms to essential, leading to an unreasonable division.

DBLP collaboration subgraph. We also perform a case study on the DB dataset (with 37,177 vertices and 131,715 edges), which is a subgraph of the DBLP co-authorship network that only contains database researchers and their co-authorships. We depict DB’s four subgraphs $C_{35} \subseteq R_{18} \subseteq C_{18}$ and R_{19} as shown in Fig. 9a. The full version of the figure, with the author represented by each vertex, is available at https://github.com/Flydragonet/Density_Decomposition_Computation.

As shown by our proposed Sandwich Theorem (Theorem 5), C_{35} , R_{18} , and C_{18} show a hierarchical structure, with their numbers of vertices and edges being $|C_{35}| = 72$, $|E(C_{35})| = 1,260$, $|R_{18}| = 131$, $|E(R_{18})| = 2,339$, $|C_{18}| = 529$, and $|E(C_{18})| = 7,398$. Compared to the original graph with $|V| = 37,177 \gg |C_{18}|$ vertices, R_{18} is tightly sandwiched between C_{35} and C_{18} . Furthermore, we illustrate the subgraph R_{19} with the dashed line in Fig. 9a. It can be observed that R_{18} and R_{19} are adjacent layers, but $C_{35} \subseteq R_{18}$ and $C_{35} \not\subseteq R_{19}$, indicating the tightness of $C_i \subseteq R_{\lfloor i/2 \rfloor}$ in the Sandwich Theorem.

¹Data source: <https://github.com/efekarakus/potter-network>, from which we extract a subgraph containing only main characters and main relationships for clarity.

²Data source: <http://vlado.fmf.uni-lj.si/pub/networks/data/DIC/TG/glossTG.htm>.

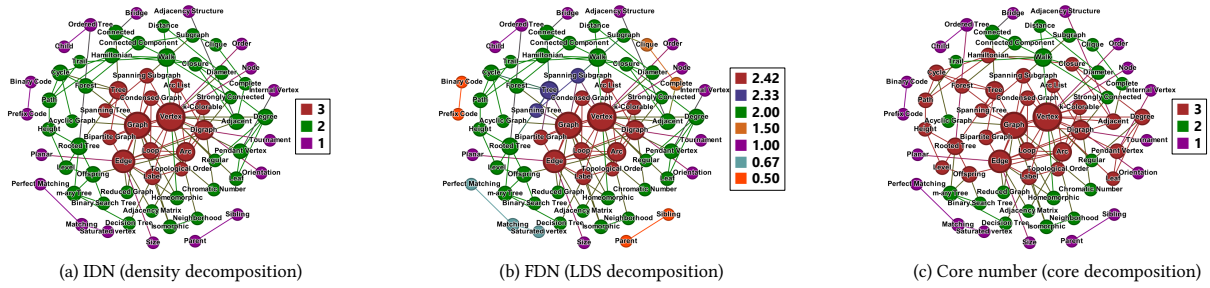


Figure 8: Dense subgraph decompositions of Graph and Digraph Glossary network GDG.

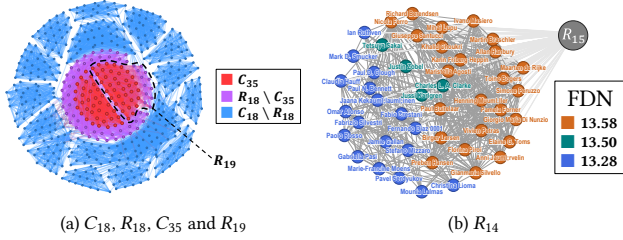


Figure 9: Dense subgraphs of decompositions of DB network.

Moreover, within DB, C_{35} and R_{19} serve as the top non-empty layer for core decomposition and density decomposition, respectively. It can be seen that C_{35} fails to separate two loosely-connected communities, whereas R_{19} contains a densely-connected community. Their respective densities are $\rho(C_{35}) = 17.5 < \rho(R_{19}) = 18.3$, and R_{19} is the densest subgraph. This again highlights the weakness of the core decomposition in capturing dense structures, particularly when contrasted with the density-aware decomposition.

To further confirm the unnecessary fine separation generated by LDS decomposition, we present the subgraph R_{14} of DB in Fig. 9b, using a single large vertex to represent R_{15} for clarity. In density decomposition, vertices in $R_{14} \setminus R_{15}$ are grouped into the same layer, while LDS decomposition divides them into three different layers with FDNs of 13.58, 13.50, and 13.28, respectively. These three layers have nearly the same outer density and form a densely connected community, suggesting that they should not be partitioned into three distinct layers. Furthermore, the four vertices with FDNs of 13.50 are connected with each other by only four edges, indicating a loose connection. It is evident that there is no need to separate these four vertices into an individual layer. All these observations suggest that LDS decomposition may overly separate the dense subgraph layers, while density decomposition, which places $R_{14} \setminus R_{15}$ in the same layer, is more reasonable.

7 RELATED WORK

Densest subgraph discovery. Our work is close to the problem of finding the densest subgraph in a graph. A well-known algorithm for computing the densest subgraph is based on invoking a parameterized Goldberg flow network $O(\log n)$ times [29]. To improve the efficiency, Danisch et al. proposed a convex programming approach to identify the densest subgraph which is considered as the state-of-the-art [23]. Several approximate algorithms have also been devised to compute the densest subgraph, such as the linear-time 2-approximation algorithm [14] $O(m+n)$ and the $(1+\epsilon)$ -approximation iterative algorithm [9, 15], with near-linear time for

each iteration. Recently, the techniques for finding the densest subgraph were also extended to the higher-order case [31, 52, 55, 58] as well as the directed graph case [38, 39].

Cohesive subgraph decomposition. Except for the density decomposition, there also exist several other cohesive subgraph decompositions aiming to decompose a graph into a set of nested cohesive subgraphs. Notable examples include the k -core decomposition [3, 36, 46], k -truss decomposition [19, 32, 57], nucleus decomposition [47–49, 51], k -edge connected subgraph decomposition [12, 13], locally densest subgraph (LDS) decomposition [43, 54], distance-generalized core decomposition [8, 22], and colorful h -star k -core decomposition [27, 28]. The core decomposition [3, 36, 46] and their variants [8, 22, 27, 28] is a degree-aware decomposition which iteratively decomposes the graph based on the degrees of the vertices. The truss decomposition is a triangle-aware decomposition that is based on the number of triangles in which an edge participates [19, 32, 57]. The nucleus decomposition can be considered as a higher-order core or truss decomposition, based on the number of cliques in which a smaller clique participates [47–49, 51]. The k -edge connected subgraph decomposition is a connectivity-aware decomposition that decomposes the graph into a set of highly connected subgraphs [12, 13]. The LDS decomposition, however, is a density-aware decomposition that can fully capture the density of the decomposed subgraphs [43, 54].

8 CONCLUSION

In this paper, we conduct an in-depth investigation of density decomposition on both static and dynamic graphs. First, we establish an inclusion relationship between density decomposition and LDS decomposition, as well as the core decomposition. For the computation of density decomposition, we develop the Flow++ algorithm using carefully-designed network flow and divide-and-conquer techniques, achieving a notable reduction in time complexity to $O(m^{3/2} \log p)$. To handle dynamic graphs, we prove a density decomposition update theorem, which shows that the maintenance of density decomposition is highly localized. Building upon this result, we develop the insertion and deletion algorithms, i.e., Insert, Delete, and Delete++, all with linear time complexity. Extensive experiments on 8 real-world datasets demonstrate the efficiency, scalability, and effectiveness of our solutions.

ACKNOWLEDGMENTS

This work was partially supported by (i) the National Key Research and Development Program of China 2021YFB3301301, (ii) NSFC-Grants U2241211 and 62072034. Rong-Hua Li is the corresponding author of this paper.

REFERENCES

- [1] Esra Akbas and Peixiang Zhao. 2017. Truss-based Community Search: a Truss-equivalence Based Indexing Approach. In *Proc. VLDB Endow.*, Vol. 10, 1298–1309.
- [2] J. Ignacio Alvarez-Hamelin, Luca Dall'Asta, Alain Barrat, and Alessandro Vespignani. 2005. Large scale networks fingerprinting and visualization using the k-core decomposition. In *NIPS*, 41–50.
- [3] Vladimir Batagelj and Matjaz Zaversnik. 2003. An $O(m)$ Algorithm for Cores Decomposition of Networks. *CoRR* cs.DS/0310049 (2003).
- [4] Jon Louis Bentley, Dorothea Haken, and James B. Saxe. 1980. A general method for solving divide-and-conquer recurrences. *SIGACT News* 12, 3 (1980), 36–44.
- [5] Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. 2013. CopyCatch: stopping group attacks by spotting lockstep behavior in social networks. In *WWW*, 119–130.
- [6] Ivona Bezáková. 2000. *Compact representations of graphs and adjacency testing*. Master's thesis. Comenius University.
- [7] Markus Blumenstock. 2016. Fast Algorithms for Pseudoarboricity. In *ALENEX*, 113–126.
- [8] Francesco Bonchi, Arijit Khan, and Lorenzo Severini. 2019. Distance-generalized Core Decomposition. In *SIGMOD*, 1006–1023.
- [9] Digvijay Boob, Yu Gao, Richard Peng, Saurabh Sawlani, Charalampos E. Tsourakakis, Di Wang, and Junxing Wang. 2020. Flowless: Extracting Densest Subgraphs Without Flow Computations. In *WWW*, 573–583.
- [10] Glencora Borradaile, Jennifer Iglesias, Theresa Migler, Antonio Ochoa, Gordon Wilfong, and Lisa Zhang. 2017. Egalitarian Graph Orientations. *Journal of Graph Algorithms and Applications* 21, 4 (2017), 687–708.
- [11] Glencora Borradaile, Theresa Migler, and Gordon T. Wilfong. 2019. Density decompositions of networks. *J. Graph Algorithms Appl.* 23, 4 (2019), 625–651.
- [12] Lijun Chang and Zhiyi Wang. 2022. A Near-Optimal Approach to Edge Connectivity-Based Hierarchical Graph Decomposition. In *Proc. VLDB Endow.*, Vol. 15, 1146–1158.
- [13] Lijun Chang, Jeffrey Xu Yu, Lu Qin, Xuemin Lin, Chengfei Liu, and Weifa Liang. 2013. Efficiently computing k-edge connected components via graph decomposition. In *SIGMOD*, 205–216.
- [14] Moses Charikar. 2000. Greedy approximation algorithms for finding dense components in a graph. In *APPROX (Lecture Notes in Computer Science, Vol. 1913)*, 84–95.
- [15] Chandra Chekuri, Kent Quanrud, and Manuel R. Torres. 2022. Densest Subgraph: Supermodularity, Iterative Peeling, and Flow. In *SODA*. SIAM, 1531–1555.
- [16] Jie Chen and Yousef Saad. 2012. Dense Subgraph Extraction with Application to Community Detection. *IEEE Trans. Knowl. Data Eng.* 24, 7 (2012), 1216–1230.
- [17] Norishige Chiba and Takao Nishizeki. 1985. Arboricity and Subgraph Listing Algorithms. *SIAM J. Comput.* 14, 1 (1985), 210–223.
- [18] Edith Cohen, Eran Halperin, Haim Kaplan, and Uri Zwick. 2003. Reachability and Distance Queries via 2-Hop Labels. *SIAM J. Comput.* 32, 5 (2003), 1338–1355.
- [19] Jonathan Cohen. 2008. Trusses: Cohesive subgraphs for social network analysis. *National security agency technical report* 16, 3.1 (2008), 1–29.
- [20] Qiangqiang Dai, Rong-Hua Li, Meihao Liao, and Guoren Wang. 2023. Maximal Defective Clique Enumeration. *Proc. ACM Manag. Data* 1, 1 (2023), 77:1–77:26.
- [21] Qiangqiang Dai, Rong-Hua Li, Hongchao Qin, Meihao Liao, and Guoren Wang. 2022. Scaling Up Maximal k-plex Enumeration. In *CIKM*, 345–354.
- [22] Qiangqiang Dai, Rong-Hua Li, Lu Qin, Guoren Wang, Weihua Yang, Zhiwei Zhang, and Ye Yuan. 2021. Scaling Up Distance-generalized Core Decomposition. In *CIKM*, 312–321.
- [23] Maximilien Danisch, T.-H. Hubert Chan, and Mauro Sozio. 2017. Large Scale Density-friendly Graph Decomposition via Convex Programming. In *WWW*, 233–242.
- [24] Luce R. Duncan and Perry Albert D. 1949. A method of matrix analysis of group structure. *Psychometrika* 14, 2 (1949), 95–116.
- [25] David Eppstein. 1994. Arboricity and Bipartite Subgraph Listing Algorithms. *Inf. Process. Lett.* 51, 4 (1994), 207–211.
- [26] Eugene Fratkin, Brian T. Naughton, Douglas L. Brutlag, and Serafim Batzoglou. 2006. MotifCut: regulatory motifs finding with maximum density subgraphs. In *Proceedings 14th International Conference on Intelligent Systems for Molecular Biology*, 156–157.
- [27] Sen Gao, Rong-Hua Li, Hongchao Qin, Hongzhi Chen, Ye Yuan, and Guoren Wang. 2022. Colorful h-star Core Decomposition. In *ICDE*, 2588–2601.
- [28] Sen Gao, Hongchao Qin, Rong-Hua Li, and Bingsheng He. 2023. Parallel Colorful h-star Core Maintenance in Dynamic Graphs. In *Proc. VLDB Endow.*, Vol. 16, 2538–2550.
- [29] Andrew V Goldberg. 1984. *Finding a maximum density subgraph*. Technical Report. University of California Berkeley, Berkeley, CA, USA.
- [30] Elfarouk Harb, Kent Quanrud, and Chandra Chekuri. 2022. Faster and scalable algorithms for densest subgraph and decomposition. *NeurIPS* 35 (2022), 26966–26979.
- [31] Yizhang He, Kai Wang, Wenjie Zhang, Xuemin Lin, and Ying Zhang. 2023. Scaling Up k-Clique Densest Subgraph Detection. *Proc. ACM Manag. Data* 1, 1 (2023), 69:1–69:26.
- [32] Xin Huang, Hong Cheng, Lu Qin, Wentao Tian, and Jeffrey Xu Yu. 2014. Querying k-truss community in large and dynamic graphs. In *SIGMOD*, 1311–1322.
- [33] Ruoming Jin, Yang Xiang, Ning Ruan, and David Fuhry. 2009. 3-HOP: a high-compression indexing scheme for reachability query. In *SIGMOD*, 813–826.
- [34] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. 1999. Trawling the Web for Emerging Cyber-Communities. *Comput. Networks* 31, 11-16 (1999), 1481–1493.
- [35] Rong-Hua Li, Qiushuo Song, Xiaokui Xiao, Lu Qin, Guoren Wang, Jeffrey Xu Yu, and Rui Mao. 2022. I/O-Efficient Algorithms for Degeneracy Computation on Massive Networks. *IEEE Trans. Knowl. Data Eng.* 34, 7 (2022), 3335–3348.
- [36] Rong-Hua Li, Jeffrey Xu Yu, and Rui Mao. 2014. Efficient Core Maintenance in Large Dynamic Graphs. *IEEE Trans. Knowl. Data Eng.* 26, 10 (2014), 2453–2465.
- [37] Chenhao Ma, Reynold Cheng, Laks V. S. Lakshmanan, and Xiaolin Han. 2022. Finding Locally Densest Subgraphs: A Convex Programming Approach. In *Proc. VLDB Endow.*, Vol. 15, 2719–2732.
- [38] Chenhao Ma, Yixiang Fang, Reynold Cheng, Laks V. S. Lakshmanan, and Xiaolin Han. 2022. A Convex-Programming Approach for Efficient Directed Densest Subgraph Discovery. In *SIGMOD*, 845–859.
- [39] Chenhao Ma, Yixiang Fang, Reynold Cheng, Laks V. S. Lakshmanan, Wenjie Zhang, and Xuemin Lin. 2021. On Directed Densest Subgraph Discovery. *ACM Trans. Database Syst.* 46, 4 (2021), 13:1–13:45.
- [40] David W. Matula and Leland L. Beck. 1983. Smallest-Last Ordering and clustering and Graph Coloring Algorithms. *J. ACM* 30, 3 (1983), 417–427.
- [41] Robert J. Mokken et al. 1979. Cliques, clubs and clans. *Quality & Quantity* 13, 2 (1979), 161–173.
- [42] Jean-Claude Picard and Maurice Queyranne. 1982. A network flow solution to some nonlinear 0-1 programming problems, with applications to graph theory. *Networks* 12, 2 (1982), 141–159.
- [43] Lu Qin, Rong-Hua Li, Lijun Chang, and Chengqi Zhang. 2015. Locally Densest Subgraph Discovery. In *KDD*, 965–974.
- [44] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAL* 4292–4293.
- [45] Barna Saha, Allison Hoch, Samir Khuller, Louiqa Raschid, and Xiao-Ning Zhang. 2010. Dense Subgraphs with Restrictions and Applications to Gene Annotation Graphs. In *RECOMB (Lecture Notes in Computer Science, Vol. 6044)*, Bonnie Berger (Ed.), Springer, 456–472.
- [46] Ahmet Erdem Sariyüce, Bugra Gedik, Gabriela Jacques-Silva, Kun-Lung Wu, and Umit V. Çatalyürek. 2013. Streaming Algorithms for k-core Decomposition. In *Proc. VLDB Endow.*, Vol. 6, 433–444.
- [47] Ahmet Erdem Sariyüce and Ali Pinar. 2016. Fast Hierarchy Construction for Dense Subgraphs. In *Proc. VLDB Endow.*, Vol. 10, 97–108.
- [48] Ahmet Erdem Sariyüce, C. Seshadhri, and Ali Pinar. 2018. Local Algorithms for Hierarchical Dense Subgraph Discovery. In *Proc. VLDB Endow.*, Vol. 12, 43–56.
- [49] Ahmet Erdem Sariyüce, C. Seshadhri, Ali Pinar, and Ümit V. Çatalyürek. 2015. Finding the Hierarchy of Dense Subgraphs using Nucleus Decompositions. In *WWW*, 927–937.
- [50] Stephen B. Seidman and Brian L. Foster. 1978. A graph-theoretic generalization of the clique concept. *Journal of Mathematical sociology* 6, 1 (1978), 139–154.
- [51] Jessica Shi, Laxman Dhulipala, and Julian Shun. 2021. Theoretically and Practically Efficient Parallel Nucleus Decomposition. In *Proc. VLDB Endow.*, Vol. 15, 583–596.
- [52] Binta Sun, Maximilien Danisch, T.-H. Hubert Chan, and Mauro Sozio. 2020. KClust++: A Simple Algorithm for Finding k-Clique Densest Subgraphs in Large Graphs. In *Proc. VLDB Endow.*, Vol. 13, 1628–1640.
- [53] Robert Endre Tarjan. 1972. Depth-First Search and Linear Graph Algorithms. *SIAM J. Comput.* 1, 2 (1972), 146–160.
- [54] Nikolaj Tatti. 2019. Density-Friendly Graph Decomposition. *ACM Trans. Knowl. Discov. Data* 13, 5 (2019), 54:1–54:29.
- [55] Charalampos E. Tsourakakis. 2015. The K-clique Densest Subgraph Problem. In *WWW*, 1122–1132.
- [56] Venkat Venkateswaran. 2004. Minimizing maximum indegree. *Discret. Appl. Math.* 143, 1-3 (2004), 374–378.
- [57] Jia Wang and James Cheng. 2012. Truss Decomposition in Massive Networks. In *Proc. VLDB Endow.*, Vol. 5, 812–823.
- [58] Yichen Xu, Chenhao Ma, Yixiang Fang, and Zhifeng Bao. 2023. Efficient and Effective Algorithms for Generalized Densest Subgraph Discovery. *Proc. ACM Manag. Data* 1, 2 (2023), 169:1–169:27.
- [59] Fan Zhang, Ying Zhang, Lu Qin, Wenjie Zhang, and Xuemin Lin. 2017. When Engagement Meets Similarity: Efficient (k, r)-Core Computation on Social Networks. In *Proc. VLDB Endow.*, Vol. 10, 998–1009.
- [60] Feng Zhao and Anthony K. H. Tung. 2012. Large Scale Cohesive Subgraphs Discovery for Social Network Visual Analysis. In *Proc. VLDB Endow.*, Vol. 6, 85–96.