

The Wedge Picking Model: A Theoretical Analysis of Graph Evolution Caused by Triadic Closure and Algorithmic Implications

Sara Ahmadian
Google New York, Google

Shahrzad Haddadan
Brown University

Social networks have become an inseparable part of human life and processing them in an efficient manner is a top priority in the study of networks. These networks are highly dynamic and they are growing incessantly. Inspired by the concept of triadic closure, we propose a probabilistic mechanism to model the evolution of these dynamic graphs. Although triadic closure is ubiquitous in social networks and its presence helps forming communities, probabilistic models encapsulating it have not been studied adequately.

We theoretically analyze our model and show how to bound the growth rate of some characteristics of the graph, such as degree of vertices. Leveraging our theoretical results, we develop a scheduling subroutine to process modifications of the graph in batches. Our scheduling subroutine is then used to speed up the state-of-the-art algorithms with negligible loss in their approximation guarantees. We demonstrate the applicability of our method by applying it to the densest subgraph and tri-densest subgraph discovery problem. Our theoretical findings are accompanied by extensive experimental evaluations.

Keywords: triadic closure, network models

INTRODUCTION

Structural properties of social network distinguish them from other graphs. These properties include bounded degeneracy, heavy tailed degree distribution, high density of triangles. Various generative models have been suggested in the literature to study or explain these properties. The most celebrated of these generative models are the preferential attachment model¹⁰, the copying model³⁷, or Chung-Lu random graph model²⁰. Each of these models explain some property of social networks but not all of them; see¹⁸ for a comparative survey. Among all of these models, only a few have encapsulated an important property of social networks: *triadic closure*^{45, 16, 15}.

Put simply, triadic closure is the tendency of a disconnected pair of users with a large common neighborhood to form connection. German sociologist George Simmel coined the term in 1908, suggesting that if there is a strong connection between A and B and there is a strong connection between A and C, then there must exist a strong or weak connection between B and C.

Triadic closure is a prominent ingredient for formation of modularity in networks⁴² and it has been shown that most of generative models such as Barabási-Albert model do not generate modularity without additionally enforcing triadic closure¹⁵. Thus, deep mathematical understanding of models explaining triadic closure is crucial. Nonetheless, the generative models based on triadic closure have not been theoretically analyzed, and little information about their growth rate or other properties is known.

Among theoretical computer scientists, a new trend of studying social networks known as “distribution free models” has recently emerged⁴⁷. Such studies leverage on specific properties of social networks to design tractable algorithms for expensive and in particular NP-hard problems. In this context the “c-closure” property was introduced by Fox et al²⁴, and it sparked notable research^{35, 31, 34}. A c-closed graph is a graph in which any two vertices v and u having a common neighborhood larger than c are connected. Problems such as listing all maximal cliques, or computing dense or sparse subgraphs, which are NP-hard in general are tractable for c-closed graphs^{24, 33}. Distribution free models and in particular c-closure have an important limitation: they fall short in modeling dynamic graphs which are more realistic in the context of social networks.

Inspired by Fox et al and subsequent work on c-closed graphs, our goal in this paper is to introduce a probabilistic dynamic model based on triadic closure, and we analyze it theoretically with the goal of designing algorithms.

In particular, we seek answer to the following questions:

1. How does triadic closure explain the evolution of dynamic graphs?
2. Can we theoretically analyze the effect of triadic closure on the growth rate of degree of vertices of a dynamic graph?
3. Can we use such bounds to design better algorithms for dynamic graphs (or to maintain previously calculated solutions)?

We provide answers to the above question by (1) introducing the *wedge picking model*, a probabilistic model for dynamic graphs based on triadic closure (2) presenting a theoretical analysis of the wedge picking model and in particular the growth rate of its local properties such as the degree of each vertex (3) after verifying that the wedge picking model explains the evolution of an input dynamic graph and estimating the parameters of the model in a pre-processing phase, we use the above theoretical results to design a speed up subroutine which we call the *rest and run* strategy. We employ our proposed method to an important primitive of graph mining: densest subgraph discovery.

The rest and run strategy works as an scheduling subroutine determining periods of “rest” and “run” for some optimization problems, thus improving their run time. The idea is to identify pre-calculated intervals for “resting” while theoretically ensuring that with high probability the quality of solution will not degrade by much. After the “rest” phase, we read a batch of updates and start the “run” phase in which we run the original algorithm.

Related Work

Study of probabilistic models for networks goes back to 1959 when pioneers of graph algorithms, Erdős and Rényi, introduced their model. Since Erdős–Rényi model does not explain most of the properties of real networks, other models like Watts–Strogatz model⁵⁴ have been introduced. The Watts–Strogatz model captures the “six degree of separation” in real network: despite having low average degree, most of vertices are reachable from each other via at most six hops. Their model also captures the existence of communities.

Real world networks have been analyzed vastly with the motivation to identify properties distinguishing them from regular random networks. Such analysis can help design of generative models. For instance, Ugander et al.⁵² analyzed social graph of active Facebook users, and reported properties such as: degree distribution, path length, component size, clustering coefficient, degeneracy, degree correlation, etc. Various models have been proposed for explaining different properties of such networks (see, e.g.,¹⁸).

Only to explain degree distribution and degree growth, researchers have suggested different models all describing the “rich gets richer” phenomenon: Barabási-Albert⁵⁵ suggest the preferential attachment model which generates graphs with power law degree distribution. Pennock et al.⁴⁴ introduce graphs with a similar degree distribution: the Discrete Gaussian Exponential distribution. Amaral et al¹¹ observe exponential cut

offs in electric power-grid graph. Among all of these models, Barabási-Albert model has been used widely and up to today particularly in the context of social networks. In this model a new vertex joins the network and it gets connected to any vertex u in the original graph with probability proportional to degree of u .

Bianconi et al.¹⁵ observed that Barabási-Albert graphs do not contain modular components like real networks do. They suggested two models to incorporate triadic closure and thus to create modularity. In the first (simpler) model, a vertex u joins the network and it gets connected to a random vertex in the graph. For each of the next edges the end point is chosen through two rules: w.p. p it is added to a randomly chosen vertex, and w.p. $1 - p$ it is added to a randomly chosen vertex from neighbors of its neighbors (thus, enforcing triadic closure). In the second model, they introduce a *fitness* factor for vertices. The probability of choosing end points for new edges will be proportional to the fitness factor. Bianconi et al's model generalizes Barabási-Albert's as setting $p = 0$ and choosing fitness factor equal to degree makes the two models equivalent.

A similar line of research uses triadic closure for link prediction: Huang et al^{30, 29} study formation of closed triads. Further they examine affects of formation of the third tie in a triad on the strength of the existing two ties²⁸. They incorporate demographics and network topology for such predictions. Estrada et al⁴ use communicability distance functions to predict triadic closure. Romero and Kleinberg⁴⁶ analyze a *link copying* regime which is similar to triadic closure. Another line of research studies estimating the number of triangles or other small patterns in a social network^{32, 36, 6, 17, 21, 7, 13} and the references within.

Despite a long list of similar work^{53, 38, 43, 40}, none of them presents a theoretical analysis of triadic closure based solely on graph topology, its affect on rate of degree growth and its possible algorithmic implications.

We finish this introduction by highlighting our paper's contribution together with a road map to read the paper.

Our Contribution

1. In Section 2, we introduce the wedge picking model which is a probabilistic generalization of c -closure, tailored to dynamic graphs. We estimate the parameters of our model in a pre-processing phase. By doing so we are also able to assert whether the evolution of input graph fits into the wedge picking model.
2. In section 3, we analyze the wedge picking model and bound degree growth (as well as tri-degree defined as the number of triangles on a vertex) for each vertex in the network.
3. In Section 4, we present our algorithmic contribution which is the introduction of the rest and run strategy. This strategy can be employed as a scheduling subroutine to boost the efficiency of other algorithms with negligible effect on their accuracy.

In Subsections 4.1 and 4.2 we show how to apply the rest and run strategy for approximating the densest subgraph and tri-densest subgraph. In particular, we incorporate the rest and run strategy in the algorithm of ²² and ⁵⁰ and we show respectively the approximation ratio of $4(1 + \epsilon)$ and $5(1 + \epsilon)$ for the modified algorithm.

4. In Section 5, we complement our theoretical results by presenting experimental analysis. Our experimental analysis is twofold: (1) We compare the quality and the run time of our algorithms in comparison to previous work. We observe that the rest and run strategy can effectively improve the speed by a factor of 1000 on large data sets while not degrading the quality of the solution. (2) We run our algorithms on two types of data sets: (a) real-world social graphs whose evolution is synthetically generated by our model, and (b) dynamic real-world graphs of "Facebook wall posts" and "YouTube users" (see Table 2). We compare the run time and accuracy of algorithms for synthetic and organic evolution and observe similar results. Thus, we demonstrate the applicability of our model on real world social graphs.

THE MODEL

In this section, we introduce our model. Let $G_t = \langle V, E_t \rangle$ be an undirected dynamic graph with a fixed vertex set V and evolving edge set E_t at time t .² The assumption of fixed vertex set is without loss of generality as we can assume that all the vertices exist in the graph from the initial state, and those vertices added to the graph in the course of evolution, initially have degree zero. Therefore, we focus on how E_t evolves through time.

Our model is a probabilistic generalization of c -closure. Thus, for any disconnected pair of vertices u and v we assume the probability that they become connected is a function of their common neighborhood denoted by $d(u, v)$. Since our experimental observations presented in Section 5 show a linear dependence, we assume this probability is $a d(u, v) + b$, for some a and b (learned in a pre-processing phase).

Edge removal occurs at a significantly lower rate compared to edge addition in social graphs. In our model, we choose uniform edge removal for deleting an edge, i.e., we pick a pair of connected vertices uniformly at random as a candidate for edge removal. Section 2.1 is devoted to a rigorous presentation of our model; this presentation has a combinatorial form which is simple and easy to analyze.

Throughout this paper, we use the following notation: Let $G_t = \langle V, E_t \rangle$ denote the graph with n vertices at time step t , G_0 denote the initial graph, and $N_t(v)$ the set of neighbors of vertex v at time t , i.e., $N_t(v) = \{u \in V \mid (u, v) \in E_t\}$. We use $N_t(u, v)$ to denote the common neighborhood of vertices u and v at time t , i.e., $N_t(u, v) = N_t(u) \cap N_t(v)$. Let $d_t(v) = |N_t(v)|$, we denote the degree of vertex v at time t and $d_t(u, v) = |N_t(u, v)|$. To denote an induced degree, induced neighborhood or induced common neighborhood in a subgraph G_0 , we mention the subgraph G_0 explicitly and use the previous notation. We use the notation $\Gamma_t(v)$ to denote the number of wedges with one end point v , i.e., the number of paths of length two starting at v , and we use $\Gamma_t(V)$ or Γ_t for short to denote the total number of wedges.

The Wedge Picking Model

Let p, q and r be constants. Throw a fair coin to decide whether to follow rule (i) or (ii):

- (i) Pick a wedge, a path of length 2, uniformly at random. If the picked wedge is open (its endpoints are disconnected), then *close* it with probability p (place an edge between the endpoints). Note that one can pick a wedge uniformly at random by picking the midpoint vertex v with probability proportional to $\binom{d_t(v)}{2}$, and then choosing the endpoints uniformly at random from $N_t(v)$. Furthermore, by picking a wedge uniformly at random the probability of placing an edge between a disconnected pair of vertices u and v will be proportional to $d_t(u, v)$.
- (ii) Pick a pair of vertices uniformly at random if they are disconnected, connect them with probability r , if they are connected, disconnect them with probability q .

So, we conclude that a disconnected pair u and v gets connected with probability $0.5(p \frac{d_t(u, v)}{\Gamma_t} + r \frac{1}{\binom{n}{2}})$.

which is linear in $d(u, v)$. Moreover, a connected pair becomes disconnected with probability $0.5(q \frac{d_t(u, v)}{\Gamma_t} + r \frac{1}{\binom{n}{2}})$.

Parameters of the Model

In this section, we explain how we learn whether or not a graph is evolving through the wedge picking model and if it does how to estimate the parameters p, q and r .

Learning p and r

In a pre-processing phase, we track edge additions between time t and $t + \Delta$ where time t is an initial time and parameter Δ will be determined later. Let $f_t(x)$ be the number of disconnected pairs with x common neighbors in G_t that become connected in $G_{t+\Delta}$, i.e., $f_t(x) = |\{(u, v) : d_t(u, v) = x, (u, v) \in E_{t+\Delta} \setminus E_t\}|$. For each $x \in \mathbb{Z}^+$, we also calculate $N(x) = |\{(u, v) : d_t(u, v) = x, (u, v) \notin E_t\}|$. Using discussion in previous section, we know that the probability of edge addition is linear in $d_t(u, v)$; we

assume it is $ad_t(u, v) + b$. We plot x vs $\frac{f(x)}{N_t(x)}$, i.e., the probability of a pair of vertices with x common neighbors becoming connected. Using linear regression we fit a line $ax + b$ to this plot and conclude:

$$\frac{f(x)}{N_t(x)} = ax + b \tag{1}$$

We use the magnitude of coefficient of determination as a certificate for guaranteeing a linear relation. For a coefficient of at least³ 0.6, we can conclude that the model follows our assumptions. Using the slope and intercept of the line then we can extract the values of p and r via Bayes's rule.

Theorem 2.1. *If a graph evolves through the wedge picking model with parameters p and r . In Equation 1, we have: $a = p \frac{M}{\Gamma_t}$ and $b = r \frac{M}{\binom{n}{2}}$, where M is the probability of an edge being added to the graph in $[t, t + \Delta]$,*

i.e., $M = \left(p \frac{\Gamma_t^o}{\Gamma_t}\right)$ where Γ_t^o is the number of open wedges at time t .

Proof. Let A be the event that an edge is added to the graph any time in $[t, t + \Delta]$. Using the definition of M and the union bound we will have $\mathbb{P}(A) \approx M \cdot |E_{t+\Delta} \setminus E_t|$. Let the B event that an edge is added to a pair with x common neighborhood. If the graph follows the wedge picking model, by Bayes rule we have:

$$\begin{aligned} \mathbb{P}(B|A) &= \frac{\mathbb{P}(A|B)\mathbb{P}(B)}{\mathbb{P}(A)} \\ &= p N_t(x) \left(\frac{x}{\Gamma_t} + r \binom{n}{2}^{-1}\right) / \mathbb{P}(A) \\ &= p N_t(d) \left(\frac{x}{\Gamma_t} + r \binom{n}{2}^{-1}\right) / M \cdot |E_{t+\Delta} \setminus E_t| \end{aligned}$$

Thus, if we take Δ so that N_t and Γ_t^o/Γ_t (subsequently M) do not change significantly (see the while condition in Algorithm 2), and using the fact that $\mathbb{P}(B|A) = \frac{f(x)}{|E_{t+\Delta} \setminus E_t|}$, we get $\frac{f(x)}{N_t(x)} = p \frac{M}{\Gamma_t} x + r \frac{M}{\binom{n}{2}}$.

Learning q . Our experimental analysis on real dynamic graphs suggests a very low rate of edge deletion compared to edge addition (see Section 5). In fact, in many of the dynamic graphs we have, edge deletion does not occur at all! Thus, instead of learning q 's value, we suffice to assert if $p/q \geq n$, which will be useful later on in the analysis of degree and tri-degree growth (see Lemma 3.1). The proof of the following theorem is straightforward and left as an exercise to the reader.

Theorem 2.2. *For a graph evolving based on the wedge picking model we have:*

$$\frac{p}{q} \geq \frac{|E_{t+\Delta} \setminus E_t|}{|E_t \setminus E_{t+\Delta}|} \cdot \frac{\Gamma_t}{\Gamma_t^o} \cdot \frac{|E_t|}{\binom{n}{2}} \tag{2}$$

THEORETICAL ANALYSIS OF DEGREE GROWTH

In real social networks, number of edges and open and closed wedges change rapidly, and keeping track of this change, or calculating them at each time step can be expensive. In this section, we show how to theoretically bound the growth rate of degrees in order to avoid expensive bookkeeping.

The rate of evolution of the graph at time t is highly dependent on the clustering coefficient of G_t (defined as the ratio of total number of open and closed wedges). When the clustering coefficient is low, we expect a high rate of evolution; meaning not only the number of edges, open and closed wedges are

changing rapidly but also the probabilities of adding or removing edges are. These rapid changes make design of algorithms a daunting task. Nevertheless, in the next lemma (Lemma 3.1) we show that these probabilities do not change significantly in limited time.

Let $\Gamma_t^k(u)$ be defined as the number of paths of length k at time t having one end at u , and $\Gamma_t^k(u,v)$ the number of paths between u and v at time t and of length k . For example, $\Gamma_t^1(u,v) = 1$ if $(u,v) \in E_t$ and 0 otherwise, $\Gamma_t^2(u,v) = d_t(u,v)$, and $\Gamma_t^1(u) = d_t(u)$. Since Γ^2 is the most commonly used notation in our paper, we use Γ to refer to Γ^2 . Note that at any time step, closing or a wedge residing on a path would decrease the length of the path. Thus, $\Gamma_t^i(u,v)$ will depend on $\Gamma_t^{i+1}(u,v), \dots, \Gamma_t^{i+\delta}(u,v)$. Moreover, for an arbitrary k and any $i < k$ by connecting two paths $\gamma_1 \in \Gamma_t^i(v)$ and $\gamma_2 \in \Gamma_t^{k-1-i}(v)$ we can create a path of length k . We use d_{max} to denote the maximum degree of the graph G_t . We formalize the above observations, in the following lemma:

Lemma 3.1. For any $u, v \in V$ and $\delta \leq e\Gamma_t / (2pd_{max}), q = o(p/d_{max})$, and $r \leq (\frac{n^2}{\Gamma_t})p$, with probability $1 - o(1)$ we have:

$$\mathbb{E}[d_{t+\delta}(u, v) - d_t(u, v)] \leq 4\delta \frac{d_t(u)d_t(v)}{\Gamma_t} \tag{3}$$


Proof. For simplicity we first assume that we have no edge deletion. We will take care of the deletions later. To warm up, consider the probability $\mathbb{P}[d_{t+1}(u, v) - d_t(u, v) > 0]$ there are two ways to add to $d_t(u, v)$: (1) by closing only one wedge placed on a path of length three between u and v , and there are two ways of doing so (rule (i)), (2) by adding an edge between $w \in N_t(v)$ and u or vice versa (rule (ii)).

Thus, $\mathbb{P}[d_{t+1}(u, v) - d_t(u, v) > 0] \leq 2\Gamma_t^3(u, v)p / \Gamma_t + (d_t(u) + d_t(v))r / n^2$, equivalently:
 $\mathbb{E}[d_{t+1}(u, v) - d_t(u, v)] \leq 2\Gamma_t^3(u, v)p / \Gamma_t + (d_t(u) + d_t(v))r / n^2$.

We now consider the general case of δ steps:

$$\mathbb{E}[d_{t+\delta}(u, v) - d_t(u, v)] = \sum_{i=1}^{\delta} \mathbb{E}[d_{t+i}(u, v) - d_{t+i-1}(u, v)] = \sum_{i=1}^{\delta} \mathbb{P}[\text{creating a wedge at } t + i]$$

FIGURE 1
THE EVOLUTION OF VERTICES' DEGREE GROWTH THROUGH TRIADIC CLOSURE



After three steps, a path of length 5 between u and v can turn to a wedge between them applying rule (i) (below picture); or alternatively two disjoint paths can turn to a wedge between u and v by applying rules (i) and (ii) (top picture). The order of edges being added is: red, blue, green.

Note that creating a wedge at time $t + j$ is possible by closing wedges (applying rule (i)) on a path which has length $j + 2$ at time t . Having rule (ii) in our hand, we can create new paths of length j at time t by placing an edge between a path of length i in $\Gamma_t^i(u)$ and a path of length $j - 1 - i$ in $\Gamma_t^{j-i-1}(v)$. Thus, the original expectation is bounded by:

$$\sum_{i=1}^{\delta} \frac{d_{t+i}(v) + d_{t+i}(u)r}{n^2} + 2\Gamma_{t+1}^3(u, v) \frac{p}{\Gamma_{t+i}} .$$

We now bound the first and second term of the sum:

$$\sum_{i=1}^{\delta} \frac{d_{t+i}(v) + d_{t+i}(u)r}{n^2} \leq \frac{\delta r(d_t(v) + d_t(u))(1 + \delta/n)}{n^2} \tag{4}$$

$$\Gamma_{t+1}^3(u, v) \leq \sum_{i=1}^{\delta} d_{max}^j \left(d_t(u)d_t(v) \frac{p^{j\kappa(2+j)}}{\Gamma_t^j} + (d_t(u) + d_t(v)) \frac{2r^j}{n^{2j}} \right) \tag{5}$$

where $\kappa(2 + j)$ is the number of ways a path of length $2 + j$ can be turned into one wedge by closing j wedges consecutively. And Γ_t is the total number of wedges at time t .

It is an easy counting problem to see $\kappa(2 + j) = (j + 2)!$. Note that for any $t, \Gamma_t \geq d_{max}^2$, and for each $j, \Gamma_t^{2+j}(u, v) \leq d_t(u)d_t(v)d_{max}^j$ and $\Gamma_t^j(u) \leq d_t(u)d_{max}^j$, where d_{max} is maximum degree at t . Note that $\frac{p^{jj!}}{d_{max}^j} \leq (p\delta/d_{max})^j$. Under the assumption that $\delta \leq e\Gamma_t/2pd_{max}$ we can simplify Equation 5 to get:

$$\sum_{i=1}^{\delta} \Gamma_{t+i}^3(u, v) \leq d_t(u)d_t(v) \sum_{i=1}^{\delta} \sum_{j=1}^i (j + 1)k^j, \quad k \leq 1/2. \tag{6}$$

Thus, as long as $r \leq (n^2/\Gamma_t)p$ (consistent with our experiments) we have:

$$\mathbb{E}[d_{t+1}(u, v) - d_t(u, v)] \leq \frac{d_t(u)d_t(v)}{\Gamma_t} (4\delta). \tag{7}$$

We now show that our initial assumption is not restrictive us as long as $q = o(\frac{p}{d_{max}})$. In fact we can show that with probability $1 - o(1)$, no edge deletion will occur in time $\delta \leq ed_{max}/p$:

$$\begin{aligned} \mathbb{P}(\text{at least one edge is deleted}) &\leq \sum_{t=1}^{\delta} q \sum_{u,v} \frac{1}{n^2} \\ \mathbb{P}(\text{at least one edge is deleted}) &\leq \delta p/\omega(d_{max}) \leq o(1) \end{aligned}$$

Having Lemma 3.1, for each $v \in V$, we can now bound the rate of its degree or tri-degree (defined later) growth.

Degree of a Vertex $u \in V$

Consider an arbitrary vertex $u \in V$, the following lemma bounds the rate of the growth of u 's degree in the wedge picking model:

Lemma 3.2. Let $\delta \leq e\Gamma_t/2pd_{max}$ and $q = o(d_{max}^{-1})p$, for any $u \in V_t$:

$$\mathbb{E}[d_{t+1}(u, v) - d_t(u, v)] \leq \delta \left(2p \frac{\Gamma_t(u)}{\Gamma_t} + r \frac{1}{n} \right) \tag{8}$$

where $\Gamma_t(u)$ is the number of wedges having one endpoint equal to v , i. e., $\Gamma_t(u) = \sum_{v \in V} d_t(u, v)$.

Proof. Let $|E_t| = m_t$, we first employ the Cauchy Schwarz inequality to show that $\Gamma_t/m_t \geq (1/2)(d_{avg} - 1)$

$$\Gamma_t = \sum_{v \in V} \frac{d_t^2(v) - d_t(v)}{2}, \quad \sum_{v \in V} d_t^2(v) \geq \frac{(\sum_{v \in V} d_t(v))^2}{n_t}; \text{ thus } 2\Gamma_t + m_t \geq m_t^2/n_t$$

$$\mathbb{E}[d_{t+1}(u, v) - d_t(u, v)] \leq \sum_{i=1}^{\delta} \mathbb{E}[d_{t+i}(u) - d_{t+i-1}(u)]$$

$$\leq \sum_{i=1}^{\delta} \sum_{u \in V} d_{t+i-1}(u, v) \frac{p}{\Gamma_t} + \frac{\delta \cdot r}{n}$$

Employing Lemma 3.1:

$$\sum_{i=1}^{\delta} \sum_{u \in V} d_{t+i-1}(u, v) \leq \sum_{i=1}^{\delta} \sum_{u \in V} d_t(u, v) + \frac{4(i-1)pd(u)d(v)}{\Gamma_t}$$

$$\leq \Gamma_t(v)\delta + \sum_{i=1}^{\delta} \frac{4(i-1)pd(u)d(v)}{\Gamma_t}$$

Having $d_t(v) \leq \Gamma_t(v)$ and $\frac{\Gamma_t}{m_t} \geq d_{avg} - 1$, we get

$$\sum_{i=1}^{\delta} \sum_{u \in V} d_{t+i-1}(u, v) \leq \left(\frac{p}{\Gamma_t}\right) \Gamma_t(v) \left(1 + \left(\frac{\delta}{2}\right) \left(\frac{4pd(v)m_t}{\Gamma_t \Gamma_t(v)}\right)\right) \delta$$

$$\leq \left(\frac{p}{\Gamma_t}\right) \Gamma_t(v) \left(1 + \left(\frac{\delta}{2}\right) \left(\frac{2p}{d_{avg} \Gamma_t}\right)\right) \delta$$

$$\leq \left(\frac{p\delta}{\Gamma_t}\right) \Gamma_t(v) \left(1 + \left(\frac{\delta}{2}\right) \left(\frac{2p}{d_{avg} \Gamma_t}\right)\right)$$

$$\leq \left(\frac{2p\delta}{\Gamma_t}\right) \Gamma_t(v) \text{ for } \delta \leq \frac{\Gamma_t}{d_{max} p}.$$

Tri-Degree of a Vertex $u \in V$

For each vertex u , and time stamp t , let tri-degree of u at time t , denoted by $D_t(u)$, be the number of triangles having u as a vertex. We have the following lemma whose proof is similar to proof of Lemma 3.2, omitted here to ease reading, and presented in the appendix:

Lemma 3.3. For any $u \in V$ and $\delta \leq ed_{max}/2p$ and $q = o(d_{max}^{-1})p$,

$$\mathbb{E}[D_{t+\delta}(u) - D_t(u)] \leq \frac{\delta \zeta_t^2(v)p}{\Gamma_t} \left(1 + \frac{8\delta p}{d_{avg}} + \frac{6\delta^2 p^2}{\Gamma_t}\right) + \frac{d_t^2(v)(\delta + \delta^2/n)}{n^2} \quad (9)$$

where $\zeta_t^2(v) \leq \sum_{u \in V} (d_t(u, v))^2$.

Remark 3.1. The growth rate of other notions of degree can be bounded similarly. These notions have been used in the literature to define various density measures [48, 23].

ALGORITHMIC IMPLICATIONS: THE REST AND RUN STRATEGY

In the dynamic graph literature, algorithms are often designed to update output solution after every single edge update; this is essential when edge updates arrive in an arbitrary (or adversarial) manner. Consider an algorithm designed for dynamic graphs which despite having low amortized cost, occasionally performs expensive operations in a row. Batching edge updates improve the run-time of these algorithms significantly. However, when not designed carefully, batching can harm the accuracy of algorithms.

In this section, we utilize our previous results to devise a speed up subroutine based on batching for certain optimization algorithms without harming their accuracy. More precisely, for a given graph G_t , we identify time Δ_t where a previously calculated solution still serves as a good approximation for any graph $G_{t'}$; $t' \in [t, t + \Delta_t]$. Our subroutine declares $[t, t + \Delta_t]$ as a “rest” interval. During rest intervals the algorithm only collects the changes in the graph and does not perform expensive operations. In the “run” intervals, the algorithm has its normal performance except since the changes are received in batches it saves redundant operations; in the experimental section we observe significant speed up due to this simple trick while maintaining the accuracy compared to the optimal solution.

Given a problem of finding a subgraph S in G satisfying a certain property, e.g., the problem of densest subgraph discovery, the question we seek to answer is:

Given an approximate solution S for G_t , what is the largest Δ_t such that with high probability, for any $t' \leq t + \Delta_t$, S is a still “good” approximation for $G_{t'}$.

Note that if an algorithm’s functionality is based on the degrees (tri-degree) of vertices of the input graph (or its degree in other induced subgraphs), Lemma 3.2 (Lemma 3.3) provides us with a value of Δ_t .

Let $C_t(v) = 2p \frac{r_t(v)}{r_t}$, and $\beta_t(1 + \varepsilon)$ be the change in degrees that the vertices in S afford to undergo while maintaining the quality within $1 + \varepsilon$ precision. We define Δ_t as follows:

$$\Delta_t := \min_{v \in V} \tau_v \quad \text{where} \quad \tau_v := \min \left\{ \frac{e r_t}{2p d_{max}}, \frac{\beta_t(1+\varepsilon)}{C_t(v)} \right\} \quad (10)$$

Employing Lemma 3.2, we conclude that using the above value for Δ_t , for any vertex u (in particular the vertices in solution S) and $\delta \in [t, t + \Delta_t]$ we will have: $\mathbb{E}[d_\delta(u) - d_t(u)] \leq (1 + \varepsilon)\beta_t$.

We formalize this observation in the next section for finding the densest subgraph in a dynamic graph. In Section 4.2 we show how to generalize these results for finding tri-densest subgraph with trivial modifications.

The Densest Subgraph Problem

In this section, to demonstrate the applicability of the rest and run strategy, we apply it to the problem of finding the densest subgraph. We present a similar algorithm for finding tri-densest subgraph (a subgraph with maximum density of triangles) in the Section 4.2. We believe that our work provides initiatives to design of algorithms for finding approximations to the *k-clique densest subgraph* problem⁴⁹.

Finding maximum cliques in a graph and even approximating it within a reasonable factor is NP hard²⁷. Thus, researchers often define density of a subgraph as its induced average degree. In the literature, *the densest subgraph problem* is discovery of subgraphs with maximum average degree (formal definition in Equation 11). This problem is polynomially tractable using flow computations²⁶, and a greedy algorithm providing a 2 approximation¹⁹, also known as *peeling* is used in practice. Peeling is a key component of most of the state-of-the-art algorithms in various massive data settings. For example: (1) in data streaming model where the network is processed in passes and the goal is to obtain a solution with minimal number of passes and low memory consumption Bahmani et al propose an algorithm based on peeling¹², (2) in the

MapReduce model where the data is distributed between different machines a number of algorithms have been suggested based on peeling^{12, 14, 22} and (3) the dynamic setting in which the graph has an evolving edge or vertex set, algorithms based on peeling have been shown to be effective^{14, 22}. A summary of the results based on peeling is presented in Table 1. It is important to mention here that peeling is not the only main strategy to find densest subgraph and in fact several foremost algorithms maintaining sketches or sparsifiers have been suggested in data streaming and dynamic graph models^{48, 23, 39, 41, 22}. We also remark that we do not aim to beat any of these algorithms theoretically. Our solution is to boost their runtime in practice by employing the *rest and run strategy*. We now formally define this problem and present our solution:

Definition: densest subgraph discovery. Let the density of graph $G = (V, E)$, denoted by $\rho(G)$, be the ratio of the number of edges in G to the number of vertices of G , i.e. $\rho(G) = |E|/|V|$. In the densest subgraph problem, our goal is to find a subgraph with maximum density. Let $\beta^*(G)$ denote the maximum density among all subgraphs of G , and let *densest*(G) be a subgraph G with maximum density, so

$$\beta^*(G) = \max_{G' \subseteq G} \rho(G'), \quad \text{densest}(G) = \operatorname{argmax}_{G' \subseteq G} \rho(G') \quad (11)$$

TABLE 1
A SUMMARY OF THE PREVIOUS WORK ON DENSEST AND TRI-DENSEST SUBGRAPH DISCOVERY BASED ON “PEELING”. MR STANDS FOR MAPREDUCE AND ACPU STANDS FOR AMORTIZED COST PER UPDATE

Authors/year	Model	Complexity	Quality	Problem
Bahmani et al, 2012 [10]	Streaming, MR	#passes: $\log_{1+\varepsilon}(n)$	$2(1 + \varepsilon)$	densest subgraph
Epasto et al, 2015 [20]	Dynamic, MR	ACPU $\log_{1+\varepsilon}(n)$	$2(1 + \varepsilon)$	densest subgraph
Bhattacharya et al, 2015 [12]	Dynamic, Streaming, MR	ACPU $\log_{1+\varepsilon}(n)$	$4(1 + \varepsilon)$	densest subgraph
Tsourakakis, 2014 [49]	Dynamic, MR	#passes: $\log_{1+\varepsilon}(n)$	$3(1 + \varepsilon)$	Tri-densest subgraph

In the dynamic graph setting algorithms of Epasto et al²² and Bhattacharya et al¹⁴, both based on the peeling strategy, present constant factor approximation algorithms, with amortized cost *poly*($\log n$) per edge addition/deletion. Despite having a low amortized cost, the above algorithms fail to output a solution in reasonable time when the size of a network (n) is huge ($n \simeq 1M$) while the edges are evolving at a high rate; see Section 5.

Our proposed solution for a faster densest subgraph discovery algorithm is to equip *peeling* with the rest and run strategy. We show that while maintaining $4(1 + \varepsilon)$ approximation for the densest subgraph problem our algorithm has significant improvement in its running time compared to its counterparts.

In this following, after briefly presenting the peeling strategy, we introduce our algorithm. In particular, our emphasis is on Epasto et al.’s algorithm²² which adopts peeling to dynamic graphs, and proves an amortized cost of $\log(n)$ per edge update, while maintaining a $2(1 + \varepsilon)$ approximate solution. We assume that the reader is familiar with their results and on a few occasions we rely on their work in our proofs (for more information see²²).

Definition: Peeling. Let β^* be the optimum density as defined in Equation 11. Partition the vertices of graph $G = (V, E)$ into layers $S_0, S_1, S_2, \dots, S_k$ as follows to find a subgraph with density at least $\beta^*/2(1 +$

ε): $S_0 = V$, while $S_{r+1} \neq \emptyset$, derive $S_{r+1} \subseteq S_r$, , by removing vertices whose degree in the induced graph S_r is less than $2\beta(1 + \varepsilon)$.

When the graph is static, a multiplicative $1 + \varepsilon$ approximation of β^* can be approximated by iterating over all $\log_{1+\varepsilon}(n)$ values. For the dynamic case Epasto et al. suggest maintaining the following invariant if possible, and if not possible rebuilding the layers by iterating over all the $\log_{1+\varepsilon}(n)$ values for β as in the static case:

Invariant: For a given density β , parameter, and current graph $G_t = \langle V_t, E \rangle$, the nodes of the graph are organized in layers $S_0, S_1, S_2, \dots, S_k$, where $S_0 = V_t$, and $S_{r+1} \subseteq S_r$ is obtained from removing vertices with degree less than $2\beta(1 + \varepsilon)$ from S_r , and $S_k = \emptyset$.

Algorithm 1 Rest and run densest sub graph.

Input: threshold c , A series of edge updates: $\{(u, v)_{t_i}, \mathfrak{N}_{t_i}\}$, where $(u, v)_{t_i}$ are pairs of vertices where the update takes place at time t_i and \mathfrak{N}_{t_i} is either addition or removal of an edge.

Output: Layers

1. At time t_0 : Run Epasto et al algorithm and find $S_0, S_1, S_2, \dots, S_k$
2. At time t_0 : Run Algorithm 2 and find p, q, r and R .
3. **if** $R \geq c$ **then**
4. **continue**
5. **repeat**
6. having p, q and r , find Δ_t
7. **while** $t_i \leq t_0 + \Delta_t$ **do**
8. Update the graph and do nothing, let U be the set of updated vertices
9. Push all the vertices in U in the stack and run algorithm of Epasto et al. to maintain $S_0, S_1, S_2, \dots, S_k$.
10. $t_i = t_i + \Delta_t$.
11. **until** ∞

Algorithm 2 Learning Phase.

Input: An initial graph G_0 , A series of edge updates: $\{(u, v)_{t_i}, \mathfrak{N}_{t_i}\}$, where $(u, v)_{t_i}$ are pairs of vertices where the update takes place at time t_i and \mathfrak{N}_{t_i} is either addition or removal of an edge.

Output: p, q, r (as defined in Section 2.1), and linear regression's coefficient of determination.

$t \leftarrow t_0$

for each $0 \leq x \leq n$, let $f(x) = 0$ and $f'(x) = 0$.

$A \leftarrow 0, D \leftarrow 0$ respectively the number of added/deleted edges.

While $\frac{\Gamma_t^o}{\Gamma_t} \in \left[\frac{1}{1+\varepsilon}, 1 + \varepsilon \right] \frac{\Gamma_0^o}{\Gamma_0}$ **and** $\forall x: N_t(x) \in \left[\frac{1}{1+\varepsilon}, (1 + \varepsilon) \right] N_0(x)$,

$t \leq |E_0|$ **do**

$x = d_t(u, v)$

Update $f(x), A, D$

Update Γ_t^o, Γ_t , and N_t

$t \leftarrow t + 1$

$a, b \leftarrow$ applying linear regression to $f(x)/N_0(x) = ax + b$.

Using Theorems 2.1 and 2.2 find p, q, r , and the coefficient of determination R .

return p, q, r, R

Our Algorithm: rest and run + Peeling. Our algorithm follows three steps: (1) We learn the parameters of the wedge picking model in a pre-processing phase (2) We find rest intervals after each run of the algorithm (3) we optimise the batch of the updated edges to feed them into the peeling algorithm.

The learning phase of our algorithm which was explained in full detail in Section 2.2 is presented as Algorithm 2. The rest and run intervals are found in the UpdatePeeling, Algorithm 1. The correctness of Algorithm 1 algorithm can be easily concluded from correctness of Epasto et al's algorithm²² and the following lemma:

Lemma 4.1. For an arbitrary vertex $v \in V$ let $d'(v)$ be degree of v in layer $S_l(v)$ at time $t + \tau(v)$. For $\beta_0 \geq c \log(n)$ for constant $c > 1$ with probability $1 - o(1)$ for all vertices v in $G \setminus S_l$, we have:

$$d'(v) \leq 4(1 + \varepsilon)\beta_0. \tag{12}$$

Proof. For each v consider a sequence of random variables $X_1, X_2, \dots, X_{\tau(v)}$, X_i indicating degree incrementation of v at time $t + i$. Note that using the values of τ defined in Equation 10 we have $\mathbb{E}[d_{t+\tau(v)}(v) - d_t(v)] \leq \beta(1 + \varepsilon)$. Applying the Chernoff bound we have: $\mathbb{P}[d_{t+\tau(v)}(v) - d_t(v) \geq 2\beta(1 + \varepsilon)] \leq e^{-(1+\varepsilon)\beta} \leq \omega\left(\frac{1}{n}\right)$.

Summing over all vertices in $G \setminus S_l$ and using the union bound we obtain:

$$\mathbb{P}[\exists v \in G \setminus S_l; d_{t+\tau(v)}(v) - d_t(v) \geq 2\beta(1 + \varepsilon)] \leq (n - \log(n))\omega\left(\frac{1}{n}\right) = o(1).$$

The following remark is concluded from Theorem [20] and Lemma 4.1:

Remark 1. Algorithm 1 maintains a $4(1 + \varepsilon)$ approximation to the densest subgraph problem. And if A is the total number of edges added to the graph in between any two consequent calls of Algorithm 1 and U is the total number of affected vertices, then since any batch of edges that are added to a single vertex in the graph are processed only once, this will cause the runtime of our algorithm to drop by a ratio of A/U compared to that of Epasto et al²².

In the next section, using Lemma 3.3 we employ the rest and run strategy to find tri-densest subgraphs (subgraphs with maximum average of tri-degree) within $5(1 + \varepsilon)$ approximation factor.

Finding Tri-Densest Subgraph

Algorithms for densest subgraph may output subgraphs with small clustering coefficient and large diameter. Hence studying other similar problems have been suggested in the literature^{50, 49, 51}: one example is searching for subgraphs which have maximal density of triangles i.e., discovery of *tri-densest subgraphs*. Note that a subgraph with a higher density of triangles has a higher clustering coefficient and is a better representative of triadic closure.

We now proceed to present our solution to this problem in the wedge picking model. Definition: Tri-densest subgraph discovery. Let the tri-density of a graph $G = (V, E)$ denoted by $\xi(G)$, be the ratio of the number of triangles in G to the number of vertices of G , i.e., $\xi(G) = |W^c(G')|/3|V'|$ where $W^c(G')$ is its number of closed wedges (Note that each triangle counts as three closed wedges). In the tri-densest subgraph problem, our goal is to find a subgraph with maximum tri-density. Let $\alpha_o(G)$ denote the maximum tri-density among all subgraphs of G , and let *tridensest*(G) be a subgraph with maximum density, so

$$\alpha_o(G) = \max_{G' \subseteq G} \xi(G'), \quad \text{tridensest}(G) = \operatorname{argmax}_{G' \subseteq G} \xi(G').$$

The problem has a $3(1 + \varepsilon)$ solution when updates are adversarial: maintain peeling layers by removing vertices with tri-degree less than $3(1+\varepsilon)\alpha$ consecutively⁵¹. Following the strategy of rest and run, at each time t we run the algorithm *UpdateTriPeeling* of⁵¹, calculate Δ_t . Then, we wait until $t + \Delta_t$ for the next call to *UpdateTriPeeling*.

Algorithm 3 Update (tri-Densest).

Input: Layers S_0, S_1, \dots, S_l, T , most recent reading time t .

Output: $5(1 + \varepsilon)$ approximation to densest subgraph, most recent reading time t^0

Let $\Delta_t = \min_v \{\tau_v \in T\}$;

Read Δ edge updates: $U = \{(u_1, v_1), (u_2, v_2), \dots, (u_{\Delta_t}, v_{\Delta_t})\}$;

$E = E \cup U$

$G_{t+\Delta_t} = \langle V, E \rangle$

Update $N = \{N(v)\}_{v \in V}$ and $T = \{\tau(v)\}_{v \in V}$ correspondingly; *UpdateTriPeeling* ($G_{t+\Delta_t}, U$); return S_l, t'

The following Theorem proves the correctness of Algorithm 4.2, and its proof is identical to proof of Lemma 4.1 and thus omitted.


Theorem 4.1. In the wedge picking model, and for $\alpha_0 \geq c \log(n)$; $c > 1$ a constant, there is an algorithm which maintains a $5(1 + \varepsilon)$ approximation of tri-densest subgraph problem with probability $1 - o(1)$ and low amortized cost.

EXPERIMENTS


Asserting the Validity of Wedge Picking Model Assumptions

In order to assert whether or not the probability of adding an edge between any pair of disconnected vertices is linear with respect to the size of their common neighbourhood, we run algorithm 2 the graphs of Table 2. Plotting $f(d)/n(d)$ over d , and running linear regression, we observe two main general behaviours for these graphs: (a) coefficient of determination is large. i.e. higher than 0.6; see e.g. Figure 2a. (b) coefficient of determination is reasonable. i.e. between 0.2 and 0.5; see e.g. Figure 2b. For most of the graphs with coefficient of determination between 0.2 and 0.5, the plot shows an interesting branching behaviour: the dependence is linear but it is in fact more that one line; see e.g. Figure 2b. In both above cases we conclude linearity.

FIGURE 2
PLOT OF $F(D)/N(D)$ OVER D AS DEFINED IN SECTION 2.2. THE GREEN LINE SHOWS THE RESULT OF LINEAR REGRESSION



(a) dataset: Digg Friends, Slope: 0.01, intercept: 0.02, coefficient of determination: 0.66. The high value of the coefficient of determination certifies linear relation.



(b) dataset: Facebook wall posts, Slope: 0.0004, intercept: .005, coefficient of determination: 0.36. Despite low coefficient of determination, we conclude linearity since we can observe two separate linear behaviors.

The Densest Subgraph Algorithm in Real Dynamic Graphs


In this section, our goal is to demonstrate the applicability of the rest and run strategy, in particular when it is applied to the densest subgraph problem. Thus, we compare the run time and accuracy of Algorithm 1 and ²² on the data sets presented in Table 2. We note that all the networks of Table 4 are not dynamic and not helpful for this section. We will later use them to generate synthetic data (Section 5.3). Besides, these datasets only evolve through edge addition, thus we set $q = 0$ and we extract p and r through learning these parameters by running Algorithm 2.

TABLE 2
WE USED THE ABOVE DYNAMIC DATASETS FROM KOBLENZ DATASET COLLECTION
IN OUR EXPERIMENTS OF SECTIONS 5.2 AND 5.3

data set	n	$ E_{tr} $	d_{avg}	d_{max}	Δ_{t_0}	Γ_0	β_{t_0}
Facebook (WOSN)	63,731	481327	25.640	1,098	1758	33044752	91.268
Digg friends	279,630	1,731,653	12.385	12,204	5276	97367436	57.34
YouTube	3,223,589	9,375,374	5.8167	91,751	39021	1245457386	33.987

Accuracy. Figure 3 shows the result of running Algorithm 1 on Digg friends dataset. We observe that algorithm tracks the density outputted by [20] closely, thus in practice the rest and run strategy does not harm accuracy of the original algorithm.

FIGURE 3
COMPARISON OF ACCURACY OF OUR ALGORITHM AGAINST [20] ON DIGG FRIENDS
DATASET. THE X AXIS IS THE NUMBER OF REST AND RUNS AND THE Y AXIS IS
THE DENSITY OUTPUT BY EACH ALGORITHM



Run time. In Table 3, we present the runtime of the algorithm for all datasets in Table 2. The runtimes demonstrate that the rest and run strategy is highly effective in reducing when employed on huge graphs. In particular, when the original algorithm does not terminate in reasonable time; see e.g., the row corresponding to YouTube.

TABLE 3
RUNTIME COMPARISON ON REAL DYNAMIC DATA SETS. THE X AXIS IS THE NUMBER OF REST AND RUNS AND THE Y AXIS IS THE DENSITY OUTPUT BY EACH ALGORITHM

data set	Our running time	[20]
Facebook (WOSN)	4s	12m
Digg Friends	20s	1m
YouTube	46s	+8h

The Densest Subgraph Algorithm in Synthetically Generated Data

In this section, our goal is to study the efficiency and accuracy of Algorithm 1 in an ideal scenario when the graph’s evolution is perfectly consistent with the assumptions of our model i.e., we have perfect linearity. To achieve this, we take $r, q = 0.001, p = 0.75$. We generate the edge sequence separately and then we run dense graph detection algorithm; doing so we achieve two goals (1) the time spent on the edge sequence generation has no affect in our algorithm runtime and (2) our comparison is based on the same benchmark for our algorithm and the algorithms we compare against.

In subsection 5.3.1 we assess the quality of approximation by comparing the density returned by our algorithm to the optimal density found by a linear program (LP) solver. We use an open source LP solver for our experiments. In Subsection 5.3.2, we switch to larger networks, having vertex set of order ten thousand to a million, to compare the quality and run time of our algorithm to ²².

Data Description. The datasets are collected from ³ and ⁵, and are listed in Table 4 and 2. We pre-process these datasets and remove duplicate edges and self loops, so for some data sets the number of edges in our input is slightly smaller than those presented in the original sources. For dynamic graphs, we sort them according to their time-stamps in our pre-processing. A more detailed description of the data sets is presented in Appendix.

Small Datasets: A Comparison of Accuracy Against Optimal Solution

We evaluate the quality of our solution by comparing it to the optimal answer. To obtain the maximum density, we solve the linear programming (LP) formulated by Charikar et al.¹⁹ In Figure 4, we present the two data sets for which we observed the largest gap between the outcomes of the two algorithms (ours and the LP). Note that the gap is always less than 1.15 after 40 rounds of rest-and-run.

Big Datasets: Runtime and Quality Assessment on Large Data Sets

We compare the efficiency of our algorithm to ²², which has amortized cost of a few microseconds per update.

FIGURE 4
QUALITY COMPARISON AGAINST OPT [17] ON DATA SETS: CA-CONDMAT (LEFT),
EMAIL-EU-CORE (RIGHT). THE X AXIS IS THE NUMBER OF REST AND RUNS
AND THE Y AXIS IS THE DENSITY OUTPUT BY EACH ALGORITHM




TABLE 4
WE USE THE ABOVE REAL WORLD GRAPHS WHICH CAN BE FOUND IN [31]

data set	n	$ Et0 $	d_{avg}	d_{max}	Δt_0	Γ_0	β_{t_0}
email-Eu-core	986	16,064	32	345	626	1,183,216	26.86
CA-GrQc	5241	14,484	5	81	147	229,867	19.56
CA-CondMat	23,133	93,439	8	279	506	1,967,650	12.5
soc-Epinions	75,879	405740	10	3044	5517	74,201,120	59.06
twitter	81,306	1,342,296	33	3383	6131	230,093,274	50.07
dblp	317,080	1,049,866	6	343	622	21,780,889	48.60
Amazon products	334,863	925,872	5	549	995	9,752,186	2.76
soc-pokec-relationships	1,632,803	22,301,964	27	14584	26919	2,086,073,558	32.68

In order to make our comparison more accurate, we feed the same sequence of updates to our algorithm and the algorithm of ²² for 100 rounds of rest and run. The only difference between our algorithm and algorithm of ²² is that we apply the changes after Δ_t updates at once to the graph and run *UpdatePeeling*. In contrast, ²² calls *UpdatePeeling* after any single update. We observe that the negligible theoretical gap between the accuracy of our algorithm vs ²² vanishes in practice, i.e., in our experiments often both solutions have the same quality. In fact, as all of our plots demonstrate, the ratio of our algorithm’s output and the output of ²² is always at most 1.65; see Figure 4.

Foremost, we observe a remarkable progress in the runtime of our algorithm. As we had expected, large values of Δ_t s boost the total running time of our algorithm to the level that it outperforms algorithms of ²² by factor of 1000. The result of our comparison is presented in Table 5.

FIGURE 5
QUALITY COMPARISON AGAINST [20] ON CA-GRQC (LEFT) AND EMAIL-EU-CORE (RIGHT) DATA SETS. THE X AXIS IS THE NUMBER OF REST AND RUNS AND THE Y AXIS IS THE DENSITY OUTPUT BY EACH ALGORITHM




TABLE 5
RUNTIME COMPARISON ON LARGER DATA SETS

data set	Our running time	[20]
Amazon products	6s	48m41s
soc-Epinions	12.4s	4h23m12s
twitter	13.52s	7h4m55s
soc-pokec-relationships	2m 18s	+24h
dblp	3.43s	32m

Data Set Description

Here we provide information on the data sets used in our experiments. This information is copied from pages on [3] describing the data set for completeness.

Email-Eu-Core Network

The network was generated using email data from a large European research institution. The data set has been anonymized. There is an edge (u, v) in the network if person u sent person v at least one email. The e-mails only represent communication between institution members (the core), and the dataset does not contain incoming messages from or outgoing messages to the rest of the world.

Collaboration Networks

We have two networks from collaboration network datasets on⁵: CA-GrQc (General Relativity and Quantum Cosmology) and CA-CondMat (General Relativity and Quantum Cosmology). These networks are from e-print arXiv on the related field. If an author i co-authored a paper with author j , the graph contains an undirected edge from i to j . If the paper is co-authored by k authors this generates a completely connected (sub)graph on k nodes.

Epinions Social Network

This is a who-trust-whom online social network of a general consumer review site Epinions.com. Members of the site can decide whether to "trust" each other. All the trust relationships interact and form the Web of Trust which is then combined with review ratings to determine which reviews are shown to the user.

Twitter

This dataset consists of 'circles' (or 'lists') from Twitter. Twitter data was crawled from public sources.

DBLP Collaboration Network

The DBLP computer science bibliography provides a comprehensive list of research papers in computer science. A co-authorship network is constructed by connecting two authors if they publish at least one paper together.

Amazon Product Co-Purchasing Network

Network was collected by crawling Amazon website. It is based on "Customers Who Bought This Item Also Bought" feature of the Amazon website. If a product i is frequently co-purchased with product j , the graph contains an undirected edge from i to j .

Pokec social network Pokec is the most popular on-line social network in Slovakia. The popularity of network has not changed even after the coming of Facebook. Pokec has been provided for more than 10 years and connects more than 1.6 million people.

ACKNOWLEDGEMENTS

The authors are thankful to Flavio Chierichetti for several helpful conversations about this paper. Shahrzad Haddadan was supported by NFS award CCF-1740741.

ENDNOTES

1. For simplicity we assume that G_t is undirected but our results can easily be modified to work for directed graphs. We defer this discussion to the full version of the paper.
2. In Algorithm 2 this threshold can be set to any arbitrary value c
3. Konect (the koblenz network collection). <http://konect.uni-koblenz.de/networks>.
4. Predicting triadic closure in networks using communicability distance functions. *SIAM J. Appl. Math.*, 75(4):1725–1744.
5. Stanford Large Network Dataset Collection. <https://snap.stanford.edu/data/>. [Online; accessed 23-May-2109].
6. S. Acer, A. Yaşar, S. Rajamanickam, M. Wolf, and V. Catalyürek. Scalable triangle counting on distributed-memory systems. In *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–5, 2019.
7. M. Agostini, M. Bressan, and S. Haddadan. Mixing time bounds for graphlet random walks. *Information Processing Letters*, page 152:105851, 2019.
8. S. Ahmadian and S. Haddadan. A theoretical analysis of graph evolution caused by triadic closure and algorithmic implications. In *2020 IEEE International Conference on Big Data (IEEE BigData 2020)*, 5-14, 2020.
9. S. Ahmadian and S. Haddadan. The wedge picking model: A dynamic graph model based on triadic closure. *CoRR*, abs/2012.01691, 2020.
10. R. A. Albert-László Barabási. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
11. S. A. B. E. M. AMARAL, L. A. N. and H. E. STANLEY. Classes of small-world networks. *Proceedings of the National Academy of Sciences*, 97, 21:11149–11152, 2000.
12. B. Bahmani, R. Kumar, and S. Vassilvitskii. Densest subgraph in streaming and mapreduce. *Proc. VLDB Endow.*, 5(5):454–465, Jan. 2012.
13. S. K. Bera and C. Seshadhri. How to count triangles, without seeing the whole graph. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery Data Mining, KDD '20*, page 306–316, New York, NY, USA, 2020. Association for Computing Machinery.
14. S. Bhattacharya, M. Henzinger, D. Nanongkai, and C. Tsourakakis. Space- and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing, STOC '15*, pages 173–182, New York, NY, USA, 2015. ACM.

15. G. Bianconi, R. K. Darst, J. Iacovacci, and S. Fortunato. Triadic closure as a basic generating mechanism of the structure of complex networks. Technical Report arXiv:1407.1664, Jul 2014. Comments: 11 pages, 14 figures.
16. J. C. Brunson. Triadic analysis of affiliation networks. *Network Science*, 3:480–508, 2015.
17. L. Bulteau, V. Froese, K. Kutzkov, and R. Pagh. Triangle counting in dynamic graph streams. In *Algorithmica* 76(1), DOI:10.1007/s00453-015-0036-4
18. D. Chakrabarti and C. Faloutsos. Graph mining: Laws, generators, and algorithms. *ACM Comput. Surv.*, 38(1), June 2006.
19. M. Charikar. Greedy approximation algorithms for finding dense components in a graph. In *Proceedings of the Third International Workshop on Approximation Algorithms for Combinatorial Optimization*, APPROX '00, pages 84–95, Berlin, Heidelberg, 2000. Springer-Verlag.
20. F. Chung and L. Lu. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics*, 6:125–145, 11 2002.
21. T. Eden, A. Levi, D. Ron, and C. Seshadhri. Approximately counting triangles in sublinear time. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 614–633, 2015.
22. A. Epasto, S. Lattanzi, and M. Sozio. Efficient densest subgraph computation in evolving graphs. In *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, pages 300–310, Republic and Canton of Geneva, Switzerland, 2015. International World Wide Web Conferences Steering Committee.
23. H. Esfandiari, M. T. Hajiaghayi, and D. P. Woodruff. Brief announcement: Applications of uniform sampling: Densest subgraph and beyond. *ArXiv*, abs/1506.04505, 2015.
24. J. Fox, T. Roughgarden, C. Seshadhri, F. Wei, and N. Wein. Finding Cliques in Social Networks: A New Distribution-Free Model. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107, pages 55:1–55:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
25. A. Gionis and C. E. Tsourakakis. Dense subgraph discovery: Kdd 2015 tutorial. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 2313–2314, New York, NY, USA, 2015. ACM.
26. A. V. Goldberg. Finding a maximum density subgraph. Technical report, Berkeley, CA, USA, 1984.
27. J. Hastad. Clique is hard to approximate within. In *Acta Mathematica*, pages 627–636, 1996.
28. H. Huang, Y. Dong, J. Tang, H. Yang, N. V. Chawla, and X. Fu. Will triadic closure strengthen ties in social networks? *ACM Trans. Knowl. Discov. Data*, 12(3), Jan. 2018.
29. H. Huang, J. Tang, L. Liu, J. Luo, and X. Fu. Triadic closure pattern analysis and prediction in social networks. *IEEE Transactions on Knowledge and Data Engineering*, 27(12):3374–3389, 2015.
30. H. Huang, J. Tang, S. Wu, L. Liu, and X. fu. Mining triadic closure patterns in social networks. *WWW '14 Companion*, page 499–504, New York, NY, USA, 2014. Association for Computing Machinery.
31. E. Husic and T. Roughgarden. Fpt algorithms for finding dense subgraphs in c -closed graphs. *CoRR*, abs/2007.09768, 2020.
32. M. Jha, C. Seshadhri, and A. Pinar. A space efficient streaming algorithm for triangle counting using the birthday paradox. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '13, page 589–597, New York, NY, USA, 2013. Association for Computing Machinery.
33. T. Koana, C. Komusiewicz, and F. Sommer. Computing dense and sparse subgraphs of weakly closed graphs, *On the 31st International Symposium on Algorithms and Computation (ISAAC 2020)*, Pages 20:1–20:17.
34. T. Koana, C. Komusiewicz, and F. Sommer. Exploiting c -closure in kernelization algorithms for graph problems, 2020. *28th Annual European Symposium on Algorithms (ESA 2020)*, pages :65:1–65:17 .
35. T. Koana and A. Nichterlein. Detecting and enumerating small induced subgraphs in c -closed graphs. *CoRR*, 2007.12077, 2020.
36. M. N. Kolountzakis, G. L. Miller, R. Peng, and C. E. Tsourakakis. Efficient triangle counting in large graphs via degree-based vertex partitioning. *International Workshop on Algorithms and Models for the Web-Graph*, WAW 2010 pages 15–24, 2010.
37. R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the web graph. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, FOCS '00, pages 57–, Washington, DC, USA, 2000. IEEE Computer Society.
38. T. Lou, J. Tang, J. Hopcroft, Z. Fang, and X. Ding. Learning to predict reciprocity and triadic closure in social networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 7, 07 2013.
39. A. McGregor, D. Tench, S. Vorotnikova, and H. T. Vu. Densest subgraph in dynamic graph streams. In *MFCS, Mathematical Foundations of Computer Science 2015*, pp 472-482 2015.

40. I. S. K. N. C. D. A. U. Milo R, Shen-Orr S. Network motifs: simple building blocks of complex networks. *Science*, 2002.
41. M. Mitzenmacher, J. W. Pachocki, R. Peng, C. E. Tsourakakis, and S. C. Xu. Scalable large near-clique detection in large-scale networks via sampling. In *KDD*, 2015.
42. M. E. J. Newman. Modularity and community structure in networks., pages 8577–8582, 2006.
43. T. Opsahl. Triadic closure in two-mode networks: Redefining the global and local clustering coefficients.
44. F. G. W. L. S. G. E. J. pennock, D. M. and C. L. Giles. Winners don't take all: Characterizing the competition for links on the web. *Proceedings of the National Academy of Sciences*, 99, 8:5207–5211, 2002.
45. G. Robins, P. Pattison, and P. Wang. Closure, connectivity and degree distributions: Exponential random graph (p^*) models for directed social networks. *Social Networks*, 31:105–117, 2009.
46. D. M. Romero and J. Kleinberg. The directed closure process in hybrid social-information networks, with an analysis of link formation on twitter. *ArXiv*, abs/1003.2469, 2010.
47. T. Roughgarden and C. Seshadhri. Beyond the worst case analysis of algorithms, Publisher: Cambridge University Press DOI: <https://doi.org/10.1017/9781108637435.035>, pp 606-625.
48. S. Sawlani and J. Wang. Near-optimal fully dynamic densest subgraph. *ArXiv*, abs/1907.03037, 2019.
49. C. Tsourakakis. The k-clique densest subgraph problem. In *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, pages 1122–1132, Republic and Canton of Geneva, Switzerland, 2015. International World Wide Web Conferences Steering Committee.
50. C. Tsourakakis, F. Bonchi, A. Gionis, F. Gullo, and M. Tsiarli. Denser than the densest subgraph: Extracting optimal quasi-cliques with quality guarantees. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '13, pages 104–112, New York, NY, USA, 2013. ACM.
51. C. E. Tsourakakis. A novel approach to finding near-cliques: The triangle-densest subgraph problem, 2014. <https://arxiv.org/pdf/1405.1477.pdf>
52. J. Ugander, B. Karrer, L. Backstrom, and C. Marlow. The anatomy of the facebook social graph. *ArXiv*, abs/1111.4503, 2011.
53. B. X. X. F. e. a. Wang, W. From triadic closure to conference closure: the role of academic conferences in promoting scientific collaborations. *Scientometrics*, 2017.
54. D. Watts and S. Strogatz. dynamics of small-world networks. volume 393:440–442, 1998.
55. S.-H. Yook, H. Jeong, and A.-L. Barabási. Modeling the internet's large-scale topology. *Proceedings of the National Academy of Sciences of the United States of America*, 99(21):13382–13386, 2002