

Estimating Clique Composition and Size Distributions from Sampled Network Data

Minas Gjoka, Emily Smith, Carter Butts

University of California, Irvine

{mgjoka, emilyjs, buttsc}@uci.edu

Abstract—Cliques are defined as complete graphs or subgraphs; they are the strongest form of cohesive subgroup, and are of interest in both social science and engineering contexts. In this paper we show how to efficiently estimate the distribution of clique sizes from a probability sample of nodes obtained from a graph (e.g., by independence or link-trace sampling). We introduce two types of unbiased estimators, one of which exploits labeling of sampled nodes neighbors and one of which does not require this information. This is the first work to present statistically principled design-based estimators for clique distributions in arbitrary graphs using sampled network data. We generalize our estimators to cases in which cliques are distinguished not only by size but also by node attributes, allowing us to estimate clique composition by size. Last, we compare our estimators on a variety of real-world graphs and provide suggestions for their use.

I. INTRODUCTION

In a large number of real-world applications it is common to represent systems, structures, or data using graphs *e.g.*, social graphs, web graphs, or protein interaction graphs. In many cases these graphs are difficult to study, most commonly because of their massive size and/or access limitations. As a result, there is a growing body of work [1,2] that uses sampling to estimate the properties of such graphs as a step towards understanding them. Furthermore, network models have been developed that receive as an input such estimated graph properties so as to generate synthetic graphs that resemble the real graph [3,4].

In this paper, we show how to efficiently estimate graph properties of the clique structure from a probability sample of nodes. Cliques are employed in a wide range of fields. In social network analysis, cliques are the foundation for studying both clustering (*i.e.*, via triangles) and cohesive subgroups [5]. Scholars have used cliques to study advice networks in entrepreneurial firms [6], informant accuracy in behavioral and cognitive network data, and friendship networks among school classes [7]. ERGM models [3] can potentially use clique distributions as an input to generate more realistic synthetic graphs. In bioinformatics, calculations of cliques have been used to identify protein structures [8] and determine optimal protein structure alignment [9]. In computer vision image recognition had been studied using maximal cliques as interest points [10].

Our estimation techniques employ an egocentric approach [5]. We first collect a uniform or non-uniform probability sample of nodes (“egos”) from the target graph. Then, we collect the egonet of each sampled node, which consists of the neighbors of the node and the edges between these neighbors. Fig. 1 presents an example that illustrates the egocentric approach. Next, we use an existing clique enumeration tool to calculate the exact clique distribution for each collected egonet. Last, we apply our unbiased estimators that combine the individual egonet calculations so as to estimate the clique distribution for the whole graph.

Our approach has several benefits. The first obvious benefit is that it allows us to estimate the clique distribution of an




Fig. 1. Illustration of the egocentric approach, clique size distribution and difference between labeled and unlabeled cases. In this case, egos 1, 6, 9 are sampled from graph G . The maximal clique size distribution of G is $C = (0, 0, 2, 2)$ since there are two maximal cliques of sizes 3 and 4.

unknown graph, as long as we have a sampling primitive that reveals the neighbors of a selected node in that graph. The second benefit is that it can be used to estimate the clique distributions of a fully known massive graph. Our approach decomposes a large problem into many smaller problems that can be independently computed, hence making the estimation parallelizable. Another benefit is the ability to estimate clique distribution in the absence of unique node labels *e.g.*, due to privacy-sensitive network data or collection limitations. Finally, our techniques can be employed with data collected using standard techniques in both online (*e.g.*, random walk or user ID sampling [11]) and offline (*e.g.*, survey instruments [12]) settings.

In summary, we make the following contributions. We present two types of unbiased estimators for the clique distribution, one that exploits labeling of nodes and one that does not require this information. To the best of our knowledge, this is the *first work to present statistically principled design-based estimators for clique distributions in arbitrary graphs using sampled network data*. We generalize the clique distribution to cases in which cliques are distinguished not only by size but also by node attributes, and provide estimators for these cases. Finally, we compare our estimators on a variety of real-world graphs and provide suggestions for their use.

II. RELATED WORK

Egocentric sampling is a widely used method for gathering network data [13,14]. This method samples individual nodes and then expands to include their neighborhoods. While this procedure does not necessarily describe the structure of the entire network, it can yield representative samples of the network [13]. Standard random sampling methods can be used to obtain egocentric network data and generalize the results to a larger population [13]. Examples of applications of egocentric sampling procedures include the network items in the General Social Survey [15] and networks obtained through crawling online social networks [11].

Although clique enumeration is in general exponential time in graph order, algorithms have been proposed for enumeration

on small and/or sparse graphs. [16] uses adjacency matrices to calculate all maximal cliques. [17] presents an algorithm to list all maximal cliques in large sparse graphs. [18] proposes an algorithm for clique enumeration with improved space complexity, which also facilitates decomposition of clique calculations on large graphs. All of these approaches are exact, but require knowledge, processing, and storage of the entire graph. In contrast, we provide for the first time estimators of clique distributions using only samples of network data.

[19] examines three methods for estimating the number of cliques in a random graph. The first is a sampling technique where they use the Bron and Kerbosch method to count a representative selection of subgraphs. The area under the curve, produced for the graph, is used to estimate the total number of cliques in the graph. The second method uses probabilistic arguments to calculate the expected number of cliques. The third is a curve fitting method. These results assume that the target graph is well-approximated by a uniform random graph. [7] outlines two approaches for estimating non-overlapping latent clique structures in a regular graph. The first uses maximum likelihood estimates and the second uses Bayesian estimates. These techniques estimate latent (not observed) cliques, and assume both disjointness and a fixed number of edges per node. These and other model-based approaches require assumptions regarding the network generating mechanism, which are avoided by our design-based approach.

Related work include methods that count the number of any subgraph (or motif) in a graph *using sampling*, such as mFinder, FANMOD [20], and MODA [21]. However, it is computationally hard for these methods to count subgraphs larger than size 10. On the other hand, our approach uses the fastest clique enumeration tool on an egonet level and estimates clique sizes up to the maximum clique size of the graph *e.g.*, up to clique size 94 for the graph soc-sign-Epinions (see Performance Evaluation Section V) and up to clique size 205 in the Facebook social graph (see our technical report [22]).

III. NOTATION AND OBJECTIVE

Let $G = (V, E)$ be an undirected graph with $N = |V|$ nodes and $|E|$ edges. For a given node $v \in V$ (“ego”), let $\mathcal{N}(v) \subseteq V$ denote the neighborhood of v in G . v 's *egonet* consists of the subgraph of G induced by v and its neighborhood (i.e., $G[v \cup \mathcal{N}(v)]$). In general, we will be concerned with settings in which we observe a probability sample of egonets from G , possibly accompanied by attributes of the associated vertices, and seek to infer properties of the clique structure of G . We elaborate as follows.

A. Clique Distributions

1) *Without Attributes*: A *clique* is an induced subgraph in which every vertex adjacent to every other vertex (i.e., a complete induced subgraph). A clique that contains i vertices is called an *order- i clique* (or *i -clique*). A clique is said to be *maximal* if no higher order clique contains it. We define C_i to be the number of order- i cliques in G , and the vector $C = (C_1, \dots, C_N)$ the *clique distribution* of G . (Counts of maximal cliques, and the maximal clique distribution are defined analogously, with the constraint that the cliques in question are maximal cliques.)

2) *With Attributes*: We generalize the clique distribution to the case in which cliques are distinguished not only by size, but also by composition. Specifically, let us assume that each vertex j in the graph has some categorical attribute X_j ; we denote the possible states of X_j by the integers $1, \dots, p$. Potential attribute distributions within cliques are defined via *composition vectors*, u , that indicate the number of vertices within a clique belonging to each of the p categories on X . Specifying u tacitly specifies order (since $\sum_{j=1}^p u_j = i$, where u_j is the number of vertices belonging to category j), and thus u can be viewed as a generalization of order for purposes of clique categorization. Accordingly, we define C_u to be the number of cliques in G that have composition vector u (*u -cliques*), with the set of all C_u all that $u \in \{u \in \mathbb{N}^p : 1 \leq \sum_{j=1}^p u_j \leq N\}$ comprising the *clique composition distribution* of G .

B. Objective

An *egocentric network sample* H_1, \dots, H_n is a probability sample of n egonets from G . Our goal is to estimate C_i and/or C_u from this sample for any given order i or composition vector u (as defined above). When estimating C_u , we assume that the egocentric network sample also contains attribute information for each ego and all neighbors.

C. Sample Properties

We assume that our egonets comprise a probability sample of the egonets in G ; that is, (i) we can treat each ego as being included in the sample with known probability, and (ii) the probability of sampling any vertex v is positive for all $v \in V$. We focus on two important cases:

Uniform Independence Sampling (UIS), where nodes are sampled independently with equal probabilities.

Non-uniform Independence Sampling (WIS), where nodes are sampled independently with probability proportional to a known weight $w(v)$.

Sampling may occur with or without replacement; we indicate these distinctions where they affect estimation. Note also that, in practice, samples drawn using link-trace methods (e.g. [11]) may closely approximate UIS or WIS, and may be employed as well. We provide an example of this approach in our technical report [22].

D. Node Labeling

When egonet H_i of ego i is sampled, it may or may not be possible to uniquely identify i 's neighbors (in the sense of knowing, e.g., whether $v \in H_i$ also belongs to some H_j). When such identification is possible, we say that the sample is *labeled*, otherwise denoting it as *unlabeled*. Fig. 1 shows the effect of labeling in an example graph in which egos 1, 6, and 9 (center graph, red) are sampled. If the sample is labeled (left), we can discern that the three sampled nodes belong to two maximal cliques: $\{1, 2, 5\}$ and $\{6, 7, 8, 9\}$. In the unlabeled case (right), however, we know only that ego 1 belongs to a maximal 3-clique, and that egos 6 and 9 each belong to one maximal 4-clique. As we will show, estimation is possible in both cases; however, labeled samples provide additional information that can be leveraged to reduce sampling error.

IV. ESTIMATION

Our goal is to infer C_i and/or C_u for some set of i and/or u from an egocentric sample. Our methods are applicable to either cliques or maximal cliques, so long as the same definition is applied consistently throughout. In what follows, we will assume that a sample of n' egos was taken from the population of vertices, of which n are unique; in the case of sampling with replacement, n may be less than n' . Our data consists of n' and n together with the sample H_1, \dots, H_n of egonets associated with unique egos (i.e., any repeatedly sampled egonets appear only once). Note that it is still possible that vertices may appear in multiple egonets, a useful fact that we exploit.

A. General Techniques

We use two basic techniques to estimate C_i and C_u , one of which can be used with either labeled or unlabeled neighborhoods and the other of which requires labeling information. We present these general techniques here, followed by specific details on how they may be used for estimation in different sampling and/or labeling scenarios.

a) Estimation via Clique Degree Sums (CDS): In this approach, we transform the problem of estimating C_i or C_u into one of estimating a sum of local structural features, and apply standard (Horvitz-Thompson) sampling theory to obtain an estimate. The approach does not require labeled neighborhoods. We begin by defining the *order- i clique degree* and *u -clique degrees* of node j , denoted d_{ij} and d_{uj} , to be respectively the numbers of i -cliques and u -cliques to which j belongs. Totaling these quantities over V gives us the corresponding *clique degree sums*,

$$D_i = \sum_{j=1}^N d_{ij} \quad \text{and} \quad D_u = \sum_{j=1}^N d_{uj}, \quad (1)$$

of which the conventional degree sum D_2 is a special case. The clique degree sums are important because of their linear relationships with C_i and C_u respectively. In particular, note that every i -clique appears i times in D_i ; since this is true irrespective of composition, we take without loss of generality $i = \sum_j u_j$ for the u -clique case to simplify exposition. From this it immediately follows that $D_i = iC_i$ and $D_u = uC_u$, and hence

$$C_i = D_i/i \quad \text{and} \quad C_u = D_u/u. \quad (2)$$

Since i is a known constant, this gives us a direct means of estimating the clique counts: for any unbiased estimators \widehat{D}_i and \widehat{D}_u of D_i and D_u , the estimators $\widehat{C}_i = \widehat{D}_i/i$ and $\widehat{C}_u = \widehat{D}_u/u$ are unbiased estimators of C_i and C_u respectively.¹ Moreover, it also follows from standard properties of the variance that $\text{Var}(\widehat{C}_i) = \text{Var}(\widehat{D}_i)/i^2$ and $\text{Var}(\widehat{C}_u) = \text{Var}(\widehat{D}_u)/u^2$ since we are merely multiplying a random variable by a constant. Thus, the sampling variability of our clique count estimators can be determined immediately from the properties of our degree sum estimators.

To estimate clique degree sums from egonet data, we employ Horvitz-Thompson (H-T) estimators [23], which are design-unbiased (with or without replacement) and which are asymptotically Gaussian for many data collection designs [24]. Define p_j as the a priori probability of H_j appearing in the sample (i.e., the j th unique ego having been sampled at any point), and let

¹An unbiased estimator, \widehat{t} of t is a statistic such that $\mathbb{E}[\widehat{t}] = t$. The result follows immediately from the linearity of expectation.

$d_i(H_j)$ and $d_u(H_j)$ be respectively the order- i clique degree and u -clique degree of ego j . The H-T estimators of D_i and D_u are then

$$\widehat{D}_i = \sum_{j=1}^n \frac{d_i(H_j)}{p_j} \quad \text{and} \quad \widehat{D}_u = \sum_{j=1}^n \frac{d_u(H_j)}{p_j}, \quad (3)$$

with the H-T estimates of C_i and C_u obtained by dividing the degree sum estimators by i . Note that since j 's ego net contains all of its cliques, d_i and d_u require only local network information; this can also greatly speed computation, as discussed in Section IV-C. Since we are concerned only with ego's clique membership, we do not require labeled neighborhoods for this method nor do our results require any assumptions regarding possible neighborhood overlap. The inclusion probabilities for sample elements arise from the design, and may differ for distinct j .

An important special case arises when sample inclusion probabilities are unequal and known only up to a constant factor (i.e., some $w_j \propto p_j$). In this case, generalized H-T estimators [24] are required for estimation of clique counts:

$$\widehat{C}_i = \frac{N \sum_{j=1}^n d_i(H_j)/w_j}{i \sum_{j=1}^n 1/w_j} \quad \text{and} \quad \widehat{C}_u = \frac{N \sum_{j=1}^n d_u(H_j)/w_j}{u \sum_{j=1}^n 1/w_j}. \quad (4)$$

The denominator in these estimators may be recognized as an H-T estimate of the total population weight (and hence a normalizer of the weights w_j). The above estimators are asymptotically unbiased [24];

b) Estimation via Distinct Clique Counting (CC): In the case of labeled neighborhoods, we can identify particular cliques across egonets. This allows us to develop H-T estimates for clique counts that do not require degree sums. Let c'_i be the number of *distinct* i -cliques in H_1, \dots, H_n that contain the egos of their respective egonets, with c'_u the corresponding number of distinct u -cliques. We define π_k to be the probability of the k th such clique appearing in the sample (i.e., the probability of selecting at least one of its members as an ego). Given π_k , H-T estimators of the clique counts are given by

$$\widehat{C}_i = \sum_{k=1}^{c'_i} \frac{1}{\pi_k} \quad \text{and} \quad \widehat{C}_u = \sum_{k=1}^{c'_u} \frac{1}{\pi_k} \quad (5)$$

Due to lack of space, estimators of the *variance* of \widehat{C}_i and \widehat{C}_u for the clique degree sums and distinct clique counting techniques are provided in our technical report [22].

B. Inclusion Probabilities

We have provided estimators for clique counts based on order and/or composition for either labeled or unlabeled egonet samples. To use them, it remains only to determine the inclusion probabilities of nodes or cliques (p and π , above). These quantities depend on the sampling design; we here provide examples for some common and important cases for sampling of OSNs in particular but also other arbitrary graphs.

Node inclusion probabilities. The simplest case for node inclusion probabilities is that in which egos are sampled uniformly at random from the population (UIS). The inclusion probabilities depend upon the total number of samples drawn ($n' \geq n$), and whether samples are drawn with or without

replacement. In the with-replacement case, an arbitrary node j fails to be selected on any given draw with probability $1 - 1/N$, and hence is ultimately included with total probability $p_j = 1 - (1 - 1/N)^{n'}$. When sampling is performed without replacement, a total of $n' = n$ of the N available nodes are drawn, any of which could be j . The resulting inclusion probability is thus simply $p_j = n'/N$.

When the probability of inclusion on any given draw is unequal, total inclusion probabilities may depend on the details of the sampling mechanism. In the common case of independent with-replacement sampling with unequal probabilities (WIS), the probability of including node j can be determined from the probability of obtaining j on any *given* draw, p'_j , by $p_j = 1 - (1 - p'_j)^{n'}$. Without-replacement inclusion probabilities with unequal are not easily summarized, but computational tools such as [25] can be employed to obtain them.

In some cases the per-draw inclusion probability may be unequal and known only up to a constant factor (i.e., $w'_j \propto p'_j$). This situation is common in e.g. random walk sampling of OSNs, where vertices are often sampled (approximately) independently with replacement, proportional to degree. In such cases, approximating p'_j by the Hansen-Hurwitz [26] estimator $p'_j \approx w'_j \left(\sum_{k=1}^{n'} 1/w'_k \right) / (n'N)$ (where the sum is over all observations, including repetitions) is a practical alternative.

Clique inclusion probabilities. The probability of sampling a clique is equal to the probability of sampling at least one of its members. In the with-replacement UIS case, this is simply $\pi_k = 1 - (1 - i/N)^{n'}$, where i is the order of the clique. When nodes are drawn uniformly without replacement, the corresponding inclusion probability becomes $\pi_k = 1 - \prod_{k=0}^{n'} \frac{N-i-k}{N-k}$ (i.e., one minus the chance of sequentially drawing n nodes from the population of non-clique members).

In the WIS case, clique inclusion probabilities will vary with the inclusion probabilities of their members. For the k th clique in an i or u -defined set, let m_k be the i -vector denoting the clique members. The resulting inclusion probability is then $\pi_k = 1 - \left(1 - \sum_{j=1}^i p'_{m_{k,j}} \right)^{n'}$, where p' contains the per-draw node sampling probabilities. More complex unequal probability designs (e.g., without replacement) do not allow simple specification of clique inclusion probabilities, but as with the nodal case these may be computed in particular cases. Since the approach to be used depends on details of the sampling procedure, we do not consider this further here.

C. Implementation considerations

The estimation of clique distributions requires the graph size N and the enumeration of cliques for each sampled egonet. In the cases when N is not known a priori, [27] provides estimators that work with sampled network data. In general, the enumeration of maximal cliques is an NP-hard problem with worst case complexity $O(3^{N/3})$. In our approach we decompose the enumeration over the whole graph to enumeration for each sampled egonet separately. Since an egonet has in the worst case size D equal to the maximum degree, the worst case complexity is $O(n * 3^{(D/3)})$ where n is the egonet sample size; this can be trivially reduced to $O(3^{(D/3)})$ via parallel computation on each egonet. It is worth noting that in real-world graphs there is often a large difference between the graph size N and maximum

degree D e.g., in the FB social graph $N = 1.11$ billion whereas $D = 5000$. Further, our method is flexible to support partial enumeration of cliques up to order- i .

The space complexity of the Clique Degree Sum method, which works only with unlabeled needs, is $O(\text{len}(C))$ which is almost negligible. On the other hand, the distinct Clique Counting methods requires $O(\sum_{i=1} C_i)$ which can be quite large depending on the graph. In Section V we describe a heuristic that helps us choose between CDS and CC.

V. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the Clique Degree Sums (CDS) and Clique Counting (CC) estimators via simulated uniform without-replacement sampling from real-world datasets. Our results shed light on the relative advantages of these estimators for assessing both order and composition distributions.

| Dataset | $ V $ | $ E $ | Average Degree | Max Degree | # Maximal Cliques | Maximum Clique Size |
|--------------------|-----------|-----------|----------------|------------|-------------------|---------------------|
| FB:New OrL. [28] | 63 392 | 816 884 | 25.77 | 1 098 | 1 538 105 | 30 |
| email-EuAll [29] | 224 832 | 339 923 | 3.02 | 7 636 | 353 194 | 16 |
| soc-sign-Epin [29] | 119 129 | 704 265 | 11.82 | 3558 | 22 219 084 | 94 |
| soc-Slashdot [29] | 77 360 | 469 179 | 12.13 | 2 539 | 823 412 | 26 |
| amazon0601 [29] | 403 364 | 2 443 309 | 12.11 | 2 752 | 1 023 558 | 11 |
| roadnet [29] | 1 087 561 | 1 541 512 | 2.83 | 9 | 1 413 059 | 4 |
| ca-CondMat [29] | 21 362 | 91 282 | 8.55 | 279 | 17 757 | 26 |
| web-Google [29] | 855 802 | 4 291 350 | 10.02 | 6 332 | 1 401 600 | 44 |
| youtube-links [29] | 1 134 894 | 2 987 623 | 5.26 | 28 754 | 3 265 955 | 17 |
| FB:UVA [30] | 17 196 | 789 321 | 91.80 | 3 182 | 4 727 791 | 42 |
| FB:UCSD [30] | 14 948 | 443 221 | 59.30 | 2 165 | 743 328 | 43 |

TABLE I: EMPIRICAL TOPOLOGIES USED IN SEC. V

A. Datasets

Table I lists the empirical networks that we use in our evaluation study. It includes several online social networks, an email communication graph, a co-authorship network, a transportation topology and a web graph; we here treat all structures as undirected graphs. For each network we list the number of nodes, number of edges, average and maximum degree, sum of maximal cliques over the whole distribution, and the maximum clique size. The numbers of nodes and sums of maximal cliques in the graphs range from thousands to millions.


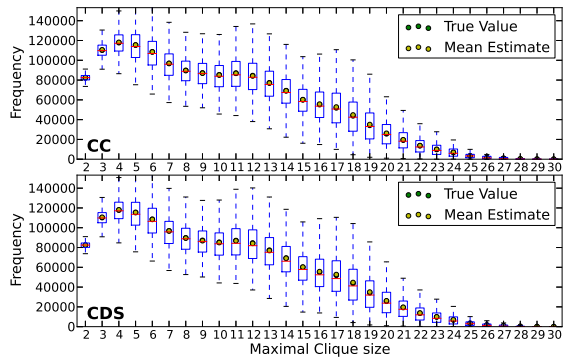
The list is composed of two groups. The first group contains no attributes and will be used to estimate C_i . The second group contains several node attributes. We have selected the node attribute “gender” to estimate C_u .

B. Error Metric

We measure the difference between estimated and actual clique distributions using the Normalized Mean Absolute Error (NMAE), defined as: $NMAE(\hat{x}, \vec{x}) = \frac{\sum (|\hat{x}_i - x_i|)}{\sum |x_i|}$, where \hat{x} and \vec{x} are the vectors that correspond to the real and estimated distributions. It returns the absolute estimation error relative to the true value, averaged over every point in the distribution.

C. Results

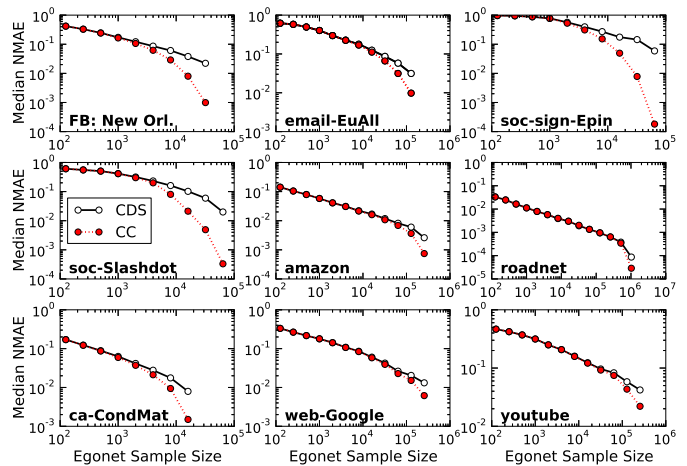

Fig. 2 shows the simulation results for the topology “FB: New Orleans” and egonet sample sizes (n) 125-32 000. We observe that with only 125 egonets the median NMAE error for both distinct CC and CDS estimators is already fairly small at ~ 0.42 . As the sample size increases, the error rapidly declines (dropping below 10%, on average, for $n \approx 4e3$). In Fig. 2(b) we “zoom in” to the case of $n = 1000$ to observe the sampling distribution of estimates for all order- i cliques in the

(a) NMAE distribution for sample sizes (n) 125–32,000(b) Distribution of simulated \widehat{C}_i versus true values; $n = 1000$, $i \in 1, \dots, 30$. Mean estimate closely matches true value due to unbiased estimation.Fig. 2. (FB: New Orleans) Clique size distribution (C_i) estimates for 1000 simulated data sets. Egonets sampled uniformly without replacement.

distribution. Confirming the unbiasedness of our estimators, the mean estimates (and in this case the median estimates) closely match the true values for all observed clique orders (1–30).

In Fig. 3 we plot the median NMAE of the CC and CDS estimators of C_i for various real-world topologies as a function of sample size. We vary n from 125 to the total size of each graph, allowing us to observe the effects of saturation on measurement error. We note that for smaller sample sizes, the CC and CDS estimators perform equally well. However, beyond a threshold sample size, the CC begins to substantially outperform the CDS estimator, reflecting the additional information associated with vertex labels. We use Table II to better interpret these results and shed some light on the causes of the “threshold” behavior. Table II contains for each topology and egonet sample size the average % of all nodes and % of all edges sampled when both egos and neighbors are included. We observe that the CC “breakaway” threshold varies for different graphs even when taking into account the total % of nodes and edges sampled. As an example, the threshold for the network “soc-sign-Epinions” is at $n \approx 4000$, corresponding to $\approx 18.1\%$ of all nodes being sampled and 34.8% of all edges being contained in some egocentric network sample on average (over 1000 simulations). On the other hand, the threshold for the network “amazon0601” is at $n = 64000$, at which point 80.9% of all nodes and 63.3% of all edges have been captured by some egocentric sample. While saturation aids the CC estimator relative to the CDS estimator, the degree of saturation required varies markedly.

Fig. 4 shows the median NMAE error of the CC and CDS estimators of C_i and C_u for several empirical networks with vertex attributes. Due to the size and density of these topologies, the egonet sample size is set between 15 – 4000. Table II shows the values for the mean % of nodes and % of edges

Fig. 3. Clique order distribution for real-world topologies. Median NMAE for the estimation of C_i calculated over 1000 simulations, as a function of n (uniform sampling without replacement).Fig. 4. Clique order and composition distributions for real-world topologies. Median NMAE for the estimation of C_i and C_u over 1000 simulations, as a function of n (uniform sampling without replacement).

sampled for these egonet sample sizes. As expected from the larger number of values (and smaller counts), estimation of C_u is at least as hard as the estimation of C_i . Depending on the composition of the attributes, the estimation of C_u ranges from being indistinguishable from C_i (see FB:UVA) or slightly worse than C_i (see FB:UCSD).

Our results show clear returns to the use of labeled neighborhoods where possible: the CC estimators perform as well or better than the CDS estimators in all cases. However, to count the distinct cliques the CC estimator needs additional space as discussed in Section IV-C. Depending on the topology, the amount of space required to implement the CC estimator might be considerably high. For example, the estimation of the clique distributions with labeling for the Facebook ’09 egocentric network samples (see our technical report [22]) requires space that is in the order of hundreds of GB. One question that naturally rises is whether there is a heuristic that can suggest whether the CC estimator is worth applying by a preliminary analysis of egonet data (*i.e.*, before we do any clique enumeration and estimation whatsoever). Two such heuristics that we examined are the use of Average Edge Count and Average Node Count as indicators of saturation. The Average Edge Count is defined and included in Table II for all networks and egonet sample sizes examined here.

Fig. 5 shows the ratio of the error between CDS and CC estimators as a function of the Average Edge Count. (Values greater than 1 favor CC.) Empirically, we observe that a value of

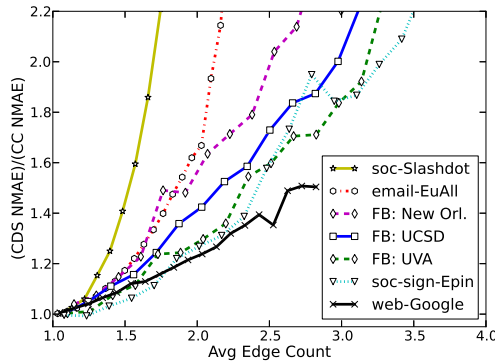


Fig. 5. Ratio of the error between CDS and CC estimators as a function of Average Edge Count i.e. a ratio of 1.30 means that the NMAE error for the CC estimator was 30% smaller than the CDC estimator. Larger values indicate greater returns to use of node labels.

| Dataset | Property | Egonet Sample Size | | | | | | | | |
|-------------------|-----------------|--------------------|-------|-------|-------|-------|-------|-------|-------|-------|
| | | 250 | 500 | 1K | 2K | 4K | 8K | 16K | 32K | 64K |
| FB:New OrL. | % nodes sampled | 9.11 | 16.16 | 26.61 | 40.04 | 55.25 | 70.56 | 84.53 | 95.38 | - |
| | % edges sampled | 5.53 | 10.55 | 19.08 | 32.12 | 49.33 | 67.82 | 84.27 | 95.75 | - |
| | avg edge count | 1.06 | 1.11 | 1.23 | 1.45 | 1.90 | 2.76 | 4.45 | 7.83 | - |
| email-EuAll | % nodes sampled | 0.38 | 0.76 | 1.37 | 2.53 | 4.65 | 8.75 | 16.35 | 30.10 | 52.66 |
| | % edges sampled | 0.47 | 0.97 | 1.86 | 3.58 | 6.70 | 12.25 | 21.61 | 36.49 | 58.66 |
| | avg edge count | 1.01 | 1.02 | 1.04 | 1.07 | 1.14 | 1.25 | 1.43 | 1.70 | 2.13 |
| soc-sign-epinions | % nodes sampled | 2.23 | 4.05 | 6.86 | 11.31 | 18.14 | 28.61 | 43.77 | 64.57 | 87.92 |
| | % edges sampled | 4.31 | 8.36 | 14.21 | 23.30 | 34.84 | 49.10 | 64.32 | 80.14 | 93.97 |
| | avg edge count | 1.07 | 1.16 | 1.33 | 1.65 | 2.19 | 3.15 | 4.78 | 7.72 | 13.10 |
| soc-slashdot | % nodes sampled | 3.60 | 6.51 | 11.15 | 18.31 | 28.71 | 43.14 | 61.75 | 82.79 | 98.83 |
| | % edges sampled | 1.68 | 3.31 | 6.20 | 11.38 | 19.90 | 33.39 | 53.15 | 77.95 | 98.43 |
| | avg edge count | 1.03 | 1.06 | 1.13 | 1.24 | 1.43 | 1.71 | 2.15 | 2.93 | 4.65 |
| amazon | % nodes sampled | 0.80 | 1.60 | 3.15 | 6.15 | 11.72 | 21.59 | 37.48 | 59.07 | 80.95 |
| | % edges sampled | 0.42 | 0.85 | 1.69 | 3.35 | 6.56 | 12.60 | 23.37 | 40.59 | 63.35 |
| | avg edge count | 1.00 | 1.01 | 1.01 | 1.02 | 1.04 | 1.08 | 1.17 | 1.35 | 1.73 |
| web-google | % nodes sampled | 0.32 | 0.62 | 1.23 | 2.39 | 4.61 | 8.69 | 15.85 | 27.40 | 43.95 |
| | % edges sampled | 0.33 | 0.65 | 1.31 | 2.57 | 5.02 | 9.61 | 17.79 | 31.05 | 49.71 |
| | avg edge count | 1.00 | 1.01 | 1.02 | 1.03 | 1.06 | 1.10 | 1.19 | 1.37 | 1.71 |
| youtube links | % nodes sampled | 0.13 | 0.26 | 0.52 | 0.98 | 1.88 | 3.51 | 6.46 | 11.65 | 20.28 |
| | % edges sampled | 0.10 | 0.21 | 0.46 | 0.85 | 1.71 | 3.27 | 6.19 | 11.48 | 20.15 |
| | avg edge count | 1.00 | 1.01 | 1.01 | 1.02 | 1.05 | 1.08 | 1.14 | 1.24 | 1.40 |
| Dataset | Property | Egonet Sample Size | | | | | | | | |
| | | 15 | 31 | 62 | 125 | 250 | 500 | 1K | 2K | 4K |
| FB: UVA | % nodes sampled | 7.42 | 14.50 | 25.66 | 41.60 | 60.11 | 75.70 | 86.58 | 92.95 | 96.61 |
| | % edges sampled | 2.20 | 4.41 | 8.47 | 15.97 | 28.28 | 44.70 | 63.90 | 81.33 | 93.07 |
| | avg edge count | 1.01 | 1.03 | 1.08 | 1.16 | 1.33 | 1.66 | 2.33 | 3.67 | 6.41 |
| FB: UCSD | % nodes sampled | 5.73 | 11.12 | 20.30 | 33.69 | 51.16 | 67.93 | 81.37 | 90.38 | 95.93 |
| | % edges sampled | 1.92 | 3.93 | 7.77 | 14.21 | 25.91 | 42.10 | 61.34 | 79.05 | 91.90 |
| | avg edge count | 1.02 | 1.03 | 1.07 | 1.13 | 1.28 | 1.59 | 2.19 | 3.38 | 5.83 |

TABLE II: UNIFORM SAMPLING WITHOUT REPLACEMENT. TOTAL % NODES SAMPLED IN THE GRAPH WHEN INCLUDING ALL EGOS AND NEIGHBORS. TOTAL % EDGES SAMPLED IN THE GRAPH WHEN INCLUDING ALL EDGES BETWEEN EGOS AND NEIGHBORS. AVERAGE EDGE COUNT: RATIO BETWEEN ALL EDGES OVER UNIQUE EDGES.

Average Edge Count above 1.5 is a heuristic indicator that a CC clique estimator for a specific egonet sample significantly outperforms the corresponding CDS clique estimator. Intuitively, the Average Edge Count of an egonet sample is correlated with the percentage of cliques that are not distinct. The higher the number of non-distinct cliques the more information the CC estimator uses compared to the CDS estimator.

Of course, in some settings (e.g., due to privacy or data collection limitations, particularly offline) is not possible to obtain information on neighbors' identities. In these cases, our simulations suggest that the CDS estimator can still provide excellent performance, even for very large graphs.

VI. CONCLUSION

In this paper, we introduced novel unbiased estimators of clique composition and size distributions based on egocentric

network samples. We presented two techniques, one of which exploits labeling of nodes (CC) and one which does not require this information (CDS). Both techniques are easily parallelizable, and suitable for use with large graphs. A Python implementation of our estimators is available at [31]. We evaluated estimator performance via simulated sampling from real-world graphs, showing that both proposed techniques work well and that CC generally outperforms CDS as the sample "saturates" the graph; because CC imposes higher space complexity, we provided a simple heuristic that suggests when gains from using CC are likely to be substantial.

REFERENCES

- [1] E. D. Kolaczyk, *Statistical Analysis of Network Data*, 2009.
- [2] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou, "Walking in Facebook: Unbiased Sampling of OSNs," in *IEEE INFOCOM '10*.
- [3] M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, and M. Morris, "statnet: Software tools for the representation, visualization, analysis and simulation of network data," *Journal of Statistical Soft.*, 2008.
- [4] M. Gjoka, M. Kurant, and A. Markopoulou, "2.5K-Graphs: from Sampling to Generation," in *Proc. of IEEE INFOCOM '13*.
- [5] S. Wasserman and K. Faust, *Social network analysis: Methods and applications*. Cambridge university press, 1994.
- [6] D. Krackhardt and M. Kilduff, "Structure, culture and simmelian ties in entrepreneurial firms," *Social networks*, vol. 24, no. 3, pp. 279–290, 2002.
- [7] I. Jansson, "Clique structure in school class data," *Social Networks*, vol. 19, no. 3, pp. 285–301, 1997.
- [8] H. M. Grindley, P. J. Artymiuk, D. W. Rice, and P. Willett, "Identification of tertiary structure resemblance in proteins using a maximal common subgraph isomorphism algorithm," *Journal of molecular biology*, 1993.
- [9] D. M. Strickland, E. Barnes, and J. S. Sokol, "Optimal protein structure alignment using maximum cliques," *Operations research*, 2005.
- [10] F. Stentiford and A. Bamidele, "Image recognition using maximal cliques of interest points," in *Proc. of IEEE ICIP '10*.
- [11] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou, "Practical recommendations on crawling online social networks," *IEEE JSAC on Measurement of Internet Topologies*, vol. 29, no. 9, pp. 1872–1892, 2011.
- [12] P. V. Marsden, "Recent developments in network measurement," *Models and methods in social network analysis*, vol. 8, p. 30, 2005.
- [13] —, "Network data and measurement," *Annual rev. of sociology*, 1990.
- [14] P. J. Carrington, J. Scott, and S. Wasserman, *Models and methods in social network analysis*. Cambridge university press, 2005.
- [15] R. S. Burt, "Network items and the general social survey," *Social networks*, vol. 6, no. 4, pp. 293–339, 1984.
- [16] E. Tomita, A. Tanaka, and H. Takahashi, "The worst-case time complexity for generating all maximal cliques and computational experiments," *Theoretical Computer Science*, vol. 363, no. 1, pp. 28–42, 2006.
- [17] D. Eppstein and D. Strash, "Listing all maximal cliques in large sparse real-world graphs," in *Experimental Algorithms*, 2011.
- [18] J. Cheng, Y. Ke, A. W.-C. Fu, J. X. Yu, and L. Zhu, "Finding maximal cliques in massive networks by h*-graph," in *Proc. of SIGMOD '10*.
- [19] S. Patel and E. Harley, "Estimation of the number of cliques in a random graph," in *Proc. of C3S2E '10*.
- [20] S. Wernicke and F. Rasche, "Fanmod: a tool for fast network motif detection," *Bioinformatics*, vol. 22, no. 9, pp. 1152–1153, 2006.
- [21] S. Omidi, F. Schreiber, and A. Masoudi-Nejad, "Moda: an efficient algorithm for network motif discovery in biological networks," *Genes & genetic systems*, vol. 84, no. 5, pp. 385–395, 2009.
- [22] M. Gjoka, E. Smith, and C. T. Butts, "Estimating Clique Composition and Size Distributions from Sampled Network Data," *arXiv:cs.SI*, 2013.
- [23] D. G. Horvitz and D. J. Thompson, "A generalization of sampling without replacement from a finite universe," *Journal of the American Statistical Association*, vol. 47, pp. 663–685, 1952.
- [24] S. K. Thompson, *Sampling*, 2nd ed. New York: Wiley, 2002.
- [25] Y. Till and A. Matei, *sampling: Survey Sampling*, 2012, r package v2.5.
- [26] M. H. Hansen and W. N. Hurwitz, "On the theory of sampling from finite populations," *The Annals of Mathematical Statistics*, vol. 14, no. 4, 1943.
- [27] L. Katzir, E. Liberty, and O. Somekh, "Estimating Sizes of Social Networks via Biased Sampling," in *Proc. of WWW*, 2011.
- [28] B. Viswanath, A. Mislove, M. Cha, and K. Gummadi, "On the evolution of user interaction in facebook," in *Proc. of WOSN '09*.
- [29] "SNAP Graph Library," <http://snap.stanford.edu/data/>.
- [30] A. Traud, P. Mucha, and M. Porter, "Social Structure of Facebook Networks," *Arxiv preprint arXiv:1102.2166*, 2011.
- [31] M. Gjoka, "Source code," <http://tinyurl.com/cliQUE-estimators>