

Scalable Positional Analysis for Studying Evolution of Nodes in Networks

Pratik Vinay Gupte*

Balaraman Ravindran†

Abstract

In social network analysis, the fundamental idea behind the notion of *position and role* is to discover actors who have similar structural signatures. Positional analysis of social networks involves partitioning the actors into disjoint sets using a notion of equivalence which captures the structure of relationships among actors. Classical approaches to Positional Analysis, such as Regular equivalence and Equitable Partitions, are too strict in grouping actors and often lead to trivial partitioning of actors in real world networks. An ϵ -Equitable Partition (ϵ EP) of a graph, which is similar in spirit to Stochastic Blockmodels, is a useful relaxation to the notion of structural equivalence which results in meaningful partitioning of actors. In this paper we propose and implement a *new scalable distributed* algorithm based on MapReduce methodology to find ϵ EP of a graph. Empirical studies on random power-law graphs show that our algorithm is highly scalable for sparse graphs, thereby giving us the ability to study positional analysis on very large scale networks. We also present the results of our algorithm on time evolving snapshots of the *facebook* and *flickr* social graphs. Results show the importance of positional analysis on large dynamic networks.

Keywords: ϵ -Equitable Partition, Structural Equivalence, Positional Analysis, Distributed Graph Partitioning

1 Introduction

Positional Analysis (PA) [26, 1] has been long used by sociologists to draw *equivalence classes* from network of social relationships. Finding equivalences in the underlying social relations gives us the ability to model social behaviour, which further aids drawing out the social and organizational structure prevalent in the network. In

PA, actors who have same structural correspondence to other actors in a network are said to occupy same “position”. As an example, head coaches in different football teams occupy the position *manager* by the virtue of the similar kind of relationship with players, assistant coaches, medical staff and the team management. It might happen that an individual coach at the position *manager* may or may not have interaction with other coaches at the same position. Given an organizational setting and the interaction patterns that exist amongst the individuals of this organization, we naturally tend to draw some abstraction around the structure and try to model its behaviour. For example, in our football team setting, the actors at the position *manager* can be a “Coach” to actors at the position *player* or a “Colleague” to the actors at the position *assistant coach*. Similarly, the actors at position *medical staff* can be a “Physiotherapist” or a “Doctor” to actors at the position *player*. While positional analysis is a very intuitive way of understanding interactions in networks, this hasn’t been widely studied for large networks due to the difficulty in developing tractable algorithms. In this paper we present a positional analysis approach that has good scaling behaviour. Hence, this work opens up the study of positions in large social networks.

The key element in finding positions, which aid in the meaningful interpretation of the data is the notion of *equivalence* used to partition the network. Classical methods of finding equivalence like structural equivalence [13], regular equivalence [27], automorphisms [9] and equitable partition [15] often lead to trivial partitioning of the actors in the network. This trivial partitioning of actors is primarily attributed due to either their strictness in the case of structural, automorphisms and equitable equivalence, which results in largely singleton position; or their leniency in the number of connections the actors at each position can possibly have with the actors at another position, as in the case of regular equivalence, which results in a giant equivalence class.

An ϵ -equitable partition (ϵ EP) [11] is a notion of equivalence, which has many advantages over the classical methods. ϵ EP allows a leeway of ϵ in the number of connections the actors at a same position can have with the actors at another position. The notion of ϵ EP is

*Index India, IIT Madras Research Park, Chennai, India 600 113. Email: pratik.gupte@gmail.com. This work was done by the author when he was a Research Scholar at the Department of Computer Science and Engineering, IIT Madras.

†Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India 600 036. Email: ravi@cse.iitm.ac.in.

The authors were partly funded by a grant from Ericsson Global Research. This manuscript is the pre-final version of the work which has been accepted at the workshop on Mining Networks and Graphs: A Big Data Analytic Challenge, to be held in conjunction with the SIAM Data Mining Conference in April 2014 (SDM-14).

similar in spirit to the notion of stochastic blockmodels [25], in that both approaches permit a bounded deviation from perfect equivalence among actors. In the Indian movies dataset from IMDB, authors in [11] have shown that actors who fall in the same cell of the partition, tend to have acted in similar kinds of movies. Further, the authors also show that people who belong to a same position of an ϵ EP tend to evolve similarly. In large social networks, tagging people who belong to the same position has potentially many advantages, both from business and individual perspective, such as, position based targeted promotions, ability to find anomalies, user churn prediction and personalised recommendations.

Though efficient graph partition refinement techniques and their application in finding the regular equivalence of a graph are well studied in the graph theoretic literature [3, 20], the application of these techniques for doing positional analysis of very large social graphs and networks is so far unknown. In this work, we propose a new algorithm to find the ϵ -equitable partition of a graph and focus on scaling this algorithm. We have successfully validated our algorithm with detailed studies on *facebook* social graph, highlighting the advantages of doing positional analysis on time evolving social network graphs. We present few results on a relatively large component of the *flickr* social network. Further more, the empirical scalability analysis of the proposed new algorithm shows that the algorithm is highly scalable for very large sparse graphs.

The contribution of our work has been twofold. First, we propose a new algorithm with better heuristics for finding the ϵ -equitable partition of a graph. Second, we have been able to parallelize this algorithm using MapReduce [6] methodology, which gives us the ability to study Positional Analysis on large dynamic social networks.

The rest of the paper is organized as follows. Section 2 discusses few mathematical preliminaries along with the definition of ϵ EP. We discuss the Fast and Parallel ϵ EP Algorithm and implementation in Section 3. We present the scalability analysis, evaluation methodology, dataset details and experimental results in Section 4. Section 5 concludes the paper with possible future directions.

2 Mathematical Preliminaries

DEFINITION 2.1. (*Partition of a graph*)

Given, graph $G \equiv \langle V, E \rangle$, V is the vertex set and E is the edge set. A partition π is defined as $\pi = \{C_1, C_2, \dots, C_n\}$ such that,

- $\cup C_i = V, i = 1$ to n and
- $i \neq j \Rightarrow C_i \cap C_j = \phi$

Thus, the definition of a *partition* of a graph G means that we have *non-empty* subsets of the vertex set V , such that all subset pairs are *disjoint* to each other. These subsets C_1, C_2, \dots, C_n are called *cells* or *blocks* of the partition π .

DEFINITION 2.2. (*Equitable partition*)

A partition $\pi = \{C_1, C_2, \dots, C_K\}$ on the vertex set V of graph G is said to be equitable [15] if, for all $1 \leq i, j \leq K, deg_G(u, C_j) = deg_G(v, C_j)$, for all $u, v \in C_i$,

where,

$$(2.1) \quad deg_G(v_i, C_j) = \text{sizeof}\{v_k \mid (v_i, v_k) \in E \text{ and } v_k \in C_j\}$$

The term $deg_G(v_i, C_j)$ denotes the number of vertices in cell C_j adjacent to the vertex v_i . An equitable partition can be used to define positions in a network; each cell C_j corresponds to a position and $deg_G(v_i, C_j)$ corresponds to the number of connections the actor v_i has to the position C_j .

McKay's algorithm (equitable refinement procedure [15]) for finding the equitable partition takes as input an ordered partition π on V and the graph G . The initial partition is usually a *unit* partition, (*i.e.*, all vertices belong to a single cell) of the graph G . An *active* list is used to hold the indices of all the unprocessed cells from π , and is updated in every iteration of the refinement procedure. c_a is the set of vertices from the *current active cell* of the partition π . The initial active cell c_a of a unit partition is therefore the entire vertex set V . Additionally, a function f , which maps every vertex $u \in V$ to its degree to c_a is used. Mathematically, $f : V \rightarrow \mathcal{N}$ is defined as follows:

$$(2.2) \quad f(u) = deg_G(u, c_a) \quad \forall u \in V$$

The procedure then sorts (in ascending order) the vertices in each cell of the partition using the value assigned to each vertex by the function f as a key for comparison. The procedure then *splits* the contents of each cell wherever the keys differ, thereby creating new cells. The partition π is updated accordingly, and the indices corresponding to any new cells formed after the split are added to the *active* list. The procedure exits when the *active* list is empty. The resulting partition is the *coarsest equitable partition* of the graph G .

DEFINITION 2.3. (*ϵ -Equitable partition*) A partition $\pi_\epsilon = \{C_1, C_2, \dots, C_K\}$ of the vertex set $\{v_1, v_2, \dots, v_n\}$, is defined as ϵ -equitable partition if, for all $1 \leq i, j \leq K, |deg_G(u, C_j) - deg_G(v, C_j)| \leq \epsilon$, for all $u, v \in C_i$,

The above definition proposes a relaxation to the strict partitioning condition of equitable partition (Definition 2.2), now equivalent actors can have a difference of ϵ in the number of connections to other cells in the partition.

DEFINITION 2.4. (*Degree vector of a vertex*)

Given a partition $\pi = \{C_1, C_2, \dots, C_K\}$ of the vertex set V of a graph G , the degree vector of a vertex $u \in V$ is defined as follows,

$$(2.3) \quad \overrightarrow{\text{deg}}_G(u) = [\text{deg}_G(u, C_1), \text{deg}_G(u, C_2), \dots, \text{deg}_G(u, C_K)]$$

Thus, the *degree vector* of a node u is a vector of size K (the total number of cells in π), where each component of the vector is the number of neighbours u has in each of the member cells of the partition π .

3 Fast and Parallel Epsilon Equitable Partition

3.1 Problems Addressed Kate and Ravindran [11] proposed an $O(n^3)$ algorithm to find the ϵ EP of a graph. We discuss this algorithm briefly*. Input to this algorithm is (i) the graph, (ii) the *coarsest equitable partition* of graph and (iii) a value of ϵ . The cells in the input equitable partition are arranged by ascending order of their *cell degrees*[†]. The algorithm then computes the *degree vector* (Equation 2.3 for each of the vertices in the graph G). The algorithm then tries to merge these cells by taking two consecutive cells at a time. If the degree vectors of the member nodes from these two cells are within ϵ distance of each other, they are merged into a single new cell. For further merging, this new cell becomes the current cell, which is then compared with the next cell for a possible merger. If the merging fails, the next cell becomes the current cell. The algorithm exits if no further merging of cells is possible. Also, the degree vectors need to be updated whenever two cells are merged. The time complexity of this algorithm to find ϵ EP of a graph is $O(n^3)$.

3.2 Fast ϵ -Equitable Partition The implementation of our Fast ϵ EP algorithm is directly based on the modification of McKay's original algorithm [15] to find the equitable partition of a graph, which iteratively refines an ordered partition until it is equitable (Section 2, Definition 2.2). The key idea in our algorithm is to allow splitting a cell only when the degrees of the member nodes of a cell are more than ϵ apart. The Fast ϵ EP algorithm and its SPLIT function is explained in Algorithm 1. The algorithm starts with the *unit* partition of the graph G and the current active cell c_a having the entire vertex set V . It then computes the function f (line 5, Algorithm 1) for each of the vertices of the graph. The algorithm then calls the SPLIT function (line 11,

Algorithm 1). The SPLIT function takes each cell from the partition π and sorts the member vertices of these cells using the function f as the comparison key (Equation 2.2, Section 2.2). Once a cell is sorted, a linear pass through the member vertices of the cell is done to check if any two consecutive vertices violate the ϵ criteria. In case of violation of the ϵ condition, the function *splits* the cell and updates the partition π and the active list accordingly. The algorithm exits either when the active list is empty or when π becomes a *discrete* partition, i.e., all cells in π are singletons.

Algorithm 1 Fast ϵ -Equitable Partition

Input: graph G , *ordered unit* partition π , epsilon ϵ

Output: ϵ -equitable partition π

```

1 active = indices( $\pi$ )
2 while (active  $\neq \phi$ ) and ( $\pi$  is not discrete) do
3   idx = min(active)
4   active = active  $\setminus \{idx\}$ 
5    $f(u) = \text{deg}_G(u, \pi[idx]) \forall u \in V$  ▷  $f : V \rightarrow \mathcal{N}$ 
6    $\pi' = \text{SPLIT}(\pi, f, \epsilon)$  ▷ line 11
7   active = active  $\cup$  [ordered indices of newly split cells
   from  $\pi'$ , while replacing (in place) the indices from  $\pi$ 
   which were split]
8    $\pi = \pi'$ 
9 end while
10 return  $\pi$ 

```

```

11 function SPLIT( $\pi, f, \epsilon$ )
12   idx = 0 ▷ index for return partition  $\pi_s$ 
13   for each currentCell in  $\pi$  do
14     sortedCell = SORT(currentCell) using  $f$  as the
   comparison key ▷ i.e., if  $f(u) < f(v)$  then  $u$  appears
   before  $v$  in sortedCell
15     currentDegree =  $f(\text{sortedCell}[0])$ 
16     for each vertex in sortedCell do
17       if ( $f(\text{vertex}) - \text{currentDegree}$ )  $\leq \epsilon$  then
18         Add vertex to cell  $\pi_s[idx]$ 
19       else
20         currentDegree =  $f(\text{vertex})$ 
21          $idx = idx + 1$ 
22         Add vertex to cell  $\pi_s[idx]$ 
23       end if
24     end for
25      $idx = idx + 1$ 
26   end for
27 return  $\pi_s$ 
28 end function

```

3.2.1 Time complexity Analysis Let n be the size of the vertex set V of G . The *while* loop of line 2 can run at most for n iterations: the case when SPLIT leads to the *discrete* partition of π , hence *active* will have n indices from $[0, 1, \dots, (n - 1)]$. The computation

*For more details on this algorithm, interested readers are kindly referred to [11] Section 3.5.

[†] *Cell or block degree* of a cell of an equitable partition is the degree of the member nodes in that cell.

of the function $f(u) = \text{deg}_G(u, c_a) \forall u \in V$ (line 5, Algorithm 1), either takes time proportional to the length of the current active cell c_a or to the length of the adjacency list of the vertex u^\ddagger . The SORT function inside SPLIT procedure (line 14, Algorithm 1) is bound to $O(n \log n)$. Also, the “splitting” (line 15 to line 24, Algorithm 1) is a linear scan and comparison of vertices in an already sorted list, hence is bound to $O(n)$. Hence, the total running time of the function SPLIT is bound to $O(n \log n)$.

The maximum cardinality of the current active cell c_a can at most be n . Further, for *dense* undirected simple graph, the maximum cardinality of the adjacency list of any vertex can also at most be $(n-1)$. Therefore for n vertices, line 5 of algorithm 1 performs in $O(n^2)$. For *sparse* graphs, the cardinality of the entire edge set is of the order of n , hence line 5 of algorithm 1 performs in the order $O(n)$.

Therefore, the total running time complexity of the proposed *Fast ϵ -Equitable Partitioning* algorithm is $O(n^3)$ for *dense* graphs and $O(n^2 \log n)$ for *sparse* graphs. In reality this would be quite less, since subsequent *splits* would only reduce the cardinality of the current *active* cell c_a . Which implies that we can safely assume that the cardinality of set c_a will be less than the cardinality of the *adjacency list* of the vertices of the graph. This analysis is only for the serial algorithm. Empirical scalability analysis on random power-law graphs shows that our parallel algorithm (Section 3.3, Algorithm 2) is an order faster in time for sparse graphs.

3.3 Parallel ϵ -Equitable Partition This section describes the parallel implementation of the Fast ϵ -Equitable Partition Algorithm 1 by MapReduce methodology [6].

In the Parallel ϵ EP Algorithm, we have implemented the most computationally intensive step of the ϵ EP algorithm, namely, computation of function f (Equation 2.2), as a MAP operation. Each *mapper* starts by initializing the *current active cell* c_a for the current iteration (line 3 Algorithm 2). The *key* input to the MAP phase is the node id n and the node data corresponding to the node n is tagged along as the *value* corresponding to this key. The MAP operation involves finding the degree of the node n to the current active cell c_a , which translates to finding the size of the intersection of the adjacency list of n and the member elements of c_a (line 5, Algorithm 2). The MAP phase *emits* the node id n as

the *key* and the degree of n to the current active cell c_a as a *value*. This corresponds to the value of function f (Equation 2.2) for the node n . Finally, a single *reducer* performs the *split* function (line 6, Algorithm 1). The output of the REDUCE phase is used by the (i) *mapper* to initialize the active cell c_a and the (ii) *reducer* itself to update the partition π and the *active* list. Single MapReduce step of the algorithm is depicted in Algorithm 2. The *iterative* MR job continues till the *active* list becomes empty or the partition becomes *discrete*.

Algorithm 2 MapReduce step of the Parallel ϵ -Equitable Partition

```

1 class MAPPER
2   method INITIALIZE()
3      $c_a \leftarrow \text{Current Active Cell}$   $\triangleright \text{active}[0]$ 
4   method MAP(id  $n$ , vertex  $N$ )
5      $d \leftarrow |N.\text{AdjacencyList} \cap c_a|$   $\triangleright d$  corresponds to
the value of function  $f(n)$ , Equation 2.2
6     EMIT(id  $n$ , value  $d$ )

```

```

1 class REDUCER  $\triangleright$  Single Reducer
2   method REDUCE()
3     SPLIT( $\pi, f, \epsilon$ )  $\triangleright$  Algorithm 1, line 11
4     UPDATE(active)  $\triangleright$  Algorithm 1, line 7
5     UPDATE( $\pi$ )

```

3.3.1 Implementation of the Parallel ϵ EP Algorithm 2 The proposed Parallel ϵ EP algorithm is *iterative* in nature, which implies that, the output of the current iteration becomes the input for the next one. The number of iterations in the Parallel ϵ EP Algorithm for sparse graphs having a million nodes is in the range of few ten thousands. The existing MapReduce framework implementations such as Hadoop and Disco [10, 19] follow the programming model and its architecture from the original MapReduce paradigm [6]. Hence, these implementations focus more on data reliability and fault tolerance guarantees in large cluster environments. This reliability and fault tolerance is usually associated with high data copy and job setup overhead costs. Although these features are suited for programs with a single MAP & a single REDUCE operation, they introduce high job setup overhead times across the *iterative* MR steps [8, 2, 29]. To circumvent this, we implemented a bare minimum MapReduce framework using open source tools GNU Parallel and rsync [22, 23]. We used GNU Parallel to trigger parallel map and reduce jobs on the cluster, rsync was used for data copy across the cluster nodes. We were able to achieve job setup overhead time in the range of **few**

[‡]Finding the *degree* of a vertex to current active cell translates to finding the *cardinality* of the *intersection set* between the current *active cell* c_a and the *adjacency list* of the vertex u . With a good choice of a data structure, the time complexity of intersection of two sets is usually proportional to the cardinality of the smaller set.

milliseconds using the custom framework, as opposed to $\sim 30 - 45$ seconds for Hadoop on a 10 node cluster isolated over a Gigabit Ethernet switch. Conceptual and detailed overview of our Lightweight MapReduce Framework implementation is depicted in Figure 1. We intelligently *sharded* the input graph data across the distributed computing nodes. The *node data partitioning* is performed based on the number of nodes n in the input graph and the number of cores p available for computation; the methodology is depicted in Figure 1(b). The *node partition* splits the input graph into nearly equal sized vertex groups for processing on each of the available cores, we cache the vertex data for each of these groups on the corresponding compute nodes. This is conceptually similar to the user control on *data persistence* and *data partitioning* in the Resilient Distributed Datasets in the *Spark* MapReduce framework [29, 28]; though our implementation was inspired independently of the Spark framework and realized before that. This helped us achieve locality in reading the input graph. Execution time empirical studies on random power-law graphs for the proposed Algorithm 2 are presented in Section 4.4.

4 Experimental Evaluation

In this section we present the results of our Fast ϵ EP algorithm. We briefly talk about the datasets used for evaluating our proposed algorithm. We also discuss the evaluation methodology and present our results. Finally, we do the scalability analysis of the proposed Parallel ϵ EP algorithm.

4.1 Datasets used for Dynamic Analysis We have used the **Facebook** (New Orleans regional network) online social network dataset from [24]. The dataset consists of timestamped friendship link formation information between September 26th, 2006 and January 22nd, 2009. We created three time evolving graph snapshots for the facebook network, the base network consists of all the links formed between September 26th, 2006 and June 22nd 2007. The remaining *two* graphs are created such that, the graph at evolved point of time $t + \delta(t)$ has the graph at time t , along with the new vertices and edges that were added to the graph between time t and time $t + \delta(t)$, with $\delta(t)$ being 290 days. Table 1 tabulates the dataset properties.

The second dataset that we used is the **Flickr** social network dataset from [18], which consists of a total of 104 days (November 2nd - December 3rd, 2006, and February 3rd - May 18th, 2007) of crawl data. This dataset consists of the timestamped link formation information among nodes. Since the nature of contact links in Flickr are directional in nature, we create an

undirected dataset as described next. For each outgoing link from user $a \rightarrow b$, if user b reciprocates the link $b \rightarrow a$, we create an undirected edge (a, b) . The time of link reciprocation by b is marked as the timestamp of the link formation. Further, we create a time evolving snapshot from this graph. The base graph G_1 has data from the first crawl, *i.e.*, between Nov 2nd - Dec 3rd, 2006. The second graph is created in a similar fashion as the **Facebook** graphs, with G_2 being G_1 plus the augmented data from the second crawl, *i.e.*, between Feb 3rd - May 18th, 2007. Table 2 tabulates the dataset properties.

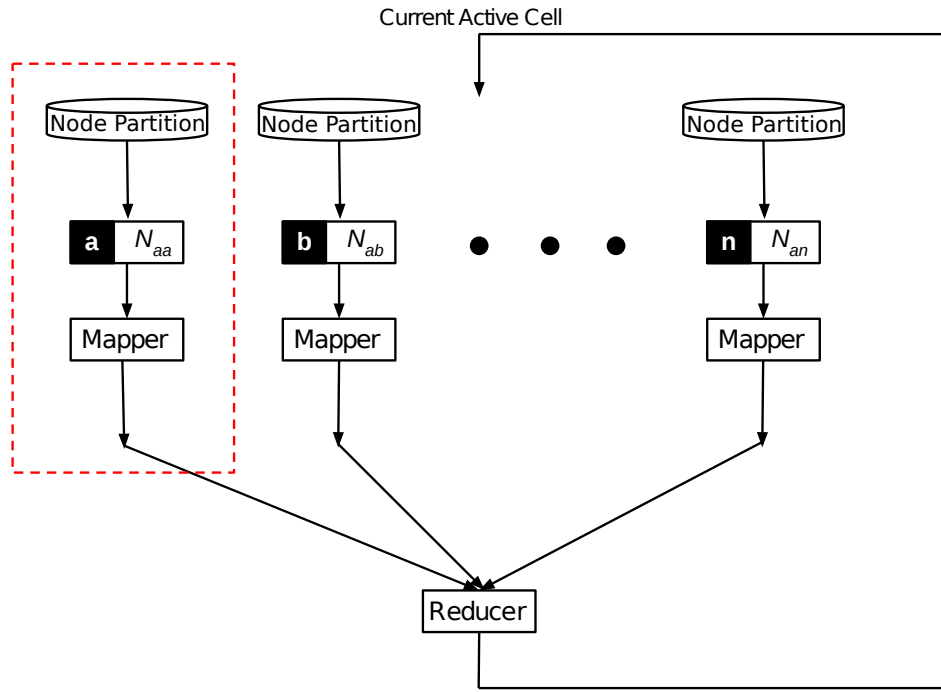
4.2 Evaluation Methodology We are primarily interested in studying the effect of PA on dynamic social networks and to characterize what role PA plays in the co-evolution of nodes in the networks. Given, a social network graph G_t at time t and its evolved network graph $G_{t+\delta t}$, our algorithm would return an ϵ -equitable partitioning π_t for G_t and $\pi_{t+\delta t}$ for $G_{t+\delta t}$. The methodology used to evaluate our proposed ϵ EP algorithm is as follows.

- i. **Partition Similarity:** We find the fraction of actors who share the same position across the partitions π_t and $\pi_{t+\delta t}$ using Equation 4.4. The new nodes in $G_{t+\delta t}$, which are not present in G_t are dropped off from $\pi_{t+\delta t}$ before computing the partition similarity score.

$$(4.4) \quad sim(\pi_t, \pi_{t+\delta t}) = \frac{1}{2} \left[\left(\frac{N - |\pi_t \cap \pi_{t+\delta t}|}{N - |\pi_t|} \right) + \left(\frac{N - |\pi_t \cap \pi_{t+\delta t}|}{N - |\pi_{t+\delta t}|} \right) \right]$$

where, N is the size of the *discrete* partition of π_t . The quantity $|\pi_t \cap \pi_{t+\delta t}|$ is the size of the partition obtained by doing *cell-wise* intersection among the cells of π_t and $\pi_{t+\delta t}$. In equation 4.4, if the number of actors who share positions across π_t and $\pi_{t+\delta t}$ is large, the value of $|\pi_t \cap \pi_{t+\delta t}|$ will be almost equal to the size of either π_t or $\pi_{t+\delta t}$. Hence, the resulting partition similarity score will be close to 1. On the other hand, if the overlap of actors between π_t and $\pi_{t+\delta t}$ is very small, $|\pi_t \cap \pi_{t+\delta t}|$ will be a large number, resulting in a similarity score close to 0. The terms in the denominator of the equation essentially provide a normalization *w.r.t.* the size of partitions π_t and $\pi_{t+\delta t}$. The value of $sim(\pi_t, \pi_{t+\delta t})$ given by Equation 4.4 is always between $[0, 1]$. We propose a fast algorithm to compute the partition similarity score in Appendix A.


- ii. **Graph theoretic network centric properties:** Given co-evolving vertex pairs (a, b) which occupy the same position in the partition π_t , we study the



(a) Conceptual Overview of our Lightweight MapReduce Implementation



$split_size = \lceil n / p \rceil$
 $start_node_index = split_size \times (partition_no - 1)$
 $end_node_index = split_size \times partition_no$
 where, n = number of nodes, p = number of cores



N_a intersection C_a



Split partition, Update Active Cell

(b) Detailed View of our MapReduce Implementation

Figure 1: (a) Conceptual overview of our Lightweight MapReduce implementation. (b) Detailed view of our MapReduce implementation. The data partition number $partition_no$ is in the range of $[1, p]$, where p is the number of available cores. The $node_partition$ splits the input graph into nearly equal sized vertex groups for processing on each of the available cores, we cache the vertex data for each of these groups on the corresponding compute nodes.

Table 1: Facebook Dataset Details

Graph	Vertices	Edges	Upto Date
1	15273	80005	2007-06-22
2	31432	218292	2008-04-07
3	61096	614796	2009-01-22

evolution of network centric properties corresponding to the vertex pairs in the time evolved graph $G_{t+\delta t}$. We study the following properties which are widely used for characterizing the graph structure:

- *Betweenness centrality* of a node v is the number of shortest paths across all the node pairs that pass through v . This signifies the importance of a node, for routing, information diffusion, etc.
- *Degree centrality* of a node is the number of nodes it is connected to in a graph. It quantifies the importance of a node *w.r.t.* the number of connections a node has.
- Counting the *number of triangles* a node is part of, is an elementary step required to find the clustering co-efficient of a node in a graph. Clustering co-efficient of a node signifies how strongly-knit a node is with its neighbourhood. There is a scalable algorithm to count the number of triangles in a graph [21].
- *Shapley value centrality* corresponds to a game theoretic notion of centrality. This models the importance of a node in information diffusion [17], it is also efficiently computable for large graphs.

We evaluate the co-evolution of nodes in various positional analysis methods using these four *network centric* properties as follows. For each pair of nodes (a, b) which occupy same position in a partition, we compute the difference $(a_t - b_t)$. Where, a_t and b_t correspond to the score of either of these four properties described previously. For the same pair of nodes, we also compute the difference $(a_{t+\delta t} - b_{t+\delta t})$. The scores $a_{t+\delta t}$ and $b_{t+\delta t}$ correspond to the property score at time $t + \delta t$. Finally, we take an absolute value of the difference of these two quantities, *i.e.*, $|(a_t - b_t) - (a_{t+\delta t} - b_{t+\delta t})|$. A low value of this quantity therefore signifies that for a co-evolving node pair (a, b) at time t , the network centric property of node a and b at time $t + \delta t$ have also evolved similarly. Note that we are not partitioning based on the centrality scores, hence our comparisons are across timestamps.

4.3 Results of Dynamic Analysis We present the evaluation of our proposed algorithm using the methodology described in the previous subsection. The results

Table 2: Flickr Dataset Details

Graph	Vertices	Edges	Upto Date
1	1277145	6042808	2006-12-03
2	1856431	10301742	2007-05-19

of the *partition similarity* score in percentages are tabulated in Table 3. We compare our method with *equitable partition* (EP) and the *degree partition*[§] (DP) for the **Facebook** dataset. We study the evolution of actors from graph $G_1 \rightarrow G_2$, $G_1 \rightarrow G_3$ and $G_2 \rightarrow G_3$, under these three partitioning methods. The ϵ -equitable partition and the degree partition display a high percentage of overlap among positions than the equitable partition. The poor performance of the EP under the partition similarity score is attributed due the strict definition of equivalence under EP. As an example, consider two nodes a and b occupy same position under EP for graph G_1 , implies that both have exactly the same degree vector. Suppose, in G_2 , the number of connections of b remained exactly the same, but node a added one extra link to another position, implies that a and b will now belong to different positions under EP. The ϵ EP consistently performs better than the DP for higher values of ϵ . The higher values of ϵ , correspond to greater bounded relaxation under the definition of ϵ EP. In most of the cases for a given graph, the number of positions under ϵ EP would decrease as we increase the ϵ . Therefore, given two ϵ EPs π_1 and π_2 , both of them would have relatively less number of positions at higher values of ϵ . This explains the higher partition similarity percentages for ϵ EP. The high values of partition similarity score for degree partition could be attributed due to the nodes in the network which don't evolve with time.

The question on choosing a correct value of ϵ , which corresponds to suitable notion of positions, while satisfying stronger cohesion among actors occupying these positions in dynamic networks is beyond the purview of this paper. Nevertheless results from Table 3 highlight a very important property of ϵ -equitable partition, namely “*tunability*”.

The study of the various *network centric properties* for co-evolving node pairs of the **Facebook** and the **Flickr** datasets, for different positional analysis methods is presented in Figure 2 and Figure 3 respectively. The *x-axis* corresponds to the bins containing the difference of a network centric property. The *y-axis* corresponds to the frequency of node pairs that belong to a particular bin, as a percentage of the total number of node pairs that occupy the same position in the partition. The results show that equitable partitioning outperforms both the ϵ EP and the DP for each of the network centric properties, which implies that they model positions of co-evolving node pairs pretty well. But the fact that equitable partition leads to trivial partition-

[§]Nodes having same degree occupy same position in the partition.




Figure 2: Co-evolving node pairs for Facebook graphs. Plots (a) to (d) $G_1 \rightarrow G_2$, Plots (e) to (h) $G_1 \rightarrow G_3$.




Figure 3: Co-evolving node pairs for Flickr graphs $G_1 \rightarrow G_2$.




Figure 4: Scalability Curve for Size of the Input vs Time, for varying Power law exponents and $\epsilon = 2, 5, 8$.

Epsilon (ϵ)	0	1	2	3	4	5	6	7	8	d*
Graph1 with Graph2	59.59	66.19	76.60	83.00	86.57	89.43	91.29	92.88	94.18	86.93
Graph1 with Graph3	54.11	57.17	69.33	76.61	80.85	84.37	86.60	88.95	90.72	79.42
Graph2 with Graph3	56.88	67.18	76.80	82.12	85.55	87.99	89.87	91.48	92.93	78.11

Table 3: Percentage of ϵ EP overlap using the Partition Similarity score (Equation 4.4) for time evolving graphs of the Facebook Network. ϵ varied from 0 to 8, $\epsilon = 0$ corresponds to an equitable partition. d* denotes the partition based on degree.

ing of nodes in a network, makes it the least suitable method for performing PA on real-world networks. Let us consider the example of the equitable partition for the graph G_3 from the **Facebook** dataset which has 61096 nodes. The EP of G_3 has 59474 cells, out of which 58494 (~96%) cells are singletons. The co-evolving node pairs under the ϵ EP outperform the DP in most of the cases for the **Facebook** networks $G_1 \rightarrow G_2$ and $G_1 \rightarrow G_3$, especially for smaller values of ϵ . The ϵ EP with smaller values for ϵ perform better because of their closeness to the equitable partition. It is worth mentioning here that, the number of positions given by ϵ EP for $\epsilon = 1$ for G_3 is 8783. This implies that, ϵ EP guarantee a high degree of confidence on the values of network centric properties of the co-evolving node pairs, along with a partitioning of a reasonable size. The **Flickr** dataset results in Figure 3 also follow a similar trend, the ϵ EP partition performs better than the DP. The percentage counts of both the properties is more spread out across initial few bins for the **Flickr** dataset. The ϵ EP for $\epsilon = 2$ has higher percentage counts for co-evolving node pairs in the bins, which correspond to smaller difference values, whereas, the co-evolving node pairs from the DP have relatively lower percentage counts in the bins closer to a difference of zero and high percentage of nodes towards the tail end of the x -axis, especially, for the shapley value centrality, which is not desirable. Also, the degree based partitions give very few positions. Therefore, ϵ EP is a consistent performer, both from the perspective of node co-evolution characteristics and number of positions it gives.

4.4 Scalability Analysis of the Parallel ϵ EP

Algorithm In this section we present empirical studies on the scalability of our proposed parallel algorithm 2. The algorithm execution was done on a single machine having eight cores; utilizing all the eight cores for the program. We study the effect of increasing the size of the input, on the running time of the algorithm. We do this analysis on *random power law* graphs by varying the power law exponent γ between $1.7 \leq \gamma \leq 2.9$. Figure 4 shows the various scalability curves. The size of the input varies from 25 thousand nodes to 1 million nodes. The running time of the algorithm increases as we decrease the value of ϵ , this is attributed due to the fact that for small values of ϵ , the number of SPLITS which we do (Algorithm 1, *line 11*) is quite large, which directly translates to increase in the number of iterations for the algorithm. Also, decreasing the power law exponent γ , increases the running time of the algorithm. Since, a lower value of γ corresponds to *denser* graphs, for dense graphs, the computation of the degree of each vertex to the current active cell c_a , therefore, becomes

a costly operation. The graph strongly suggests that the algorithm scales almost linearly with increase in the input graph size for the values of $2.1 \leq \gamma \leq 2.9$. It is worth mentioning here that, for most of the real-world graphs, γ lies between 2 and 3 [4], with few exemptions. We also performed curve-fitting using polynomial regression to get a complexity bound on the algorithm for $\epsilon = 5$. We get a running time bound of $O(n)$, $O(n \log n)$ and $O(n^2)$ for random power law graphs generated using $\gamma = 2.9$, 2.5 and 2.1 respectively. It is worth noting that, the sum squared residual for $\gamma = 2.1$ and the curves $n \log n$ and n^2 was quite marginal.

5 Conclusion and Future Work

In this paper we have presented a scalable and distributed ϵ -equitable partition algorithm. To the best of our knowledge, this is the first attempt at doing positional analysis on a large scale online social network dataset. We have been able to compute ϵ EP for a significantly large component of the **Flickr** social graph using our Parallel ϵ EP algorithm and its implementation. Further, the results of our algorithm on the **Facebook** and **Flickr** datasets show that ϵ EP is a promising tool for analyzing the evolution of nodes in dynamic networks. Empirical scalability studies on random power law graphs show that our algorithm is highly scalable for very large *sparse* graphs.

In future, it would be interesting to explore the implied advantage of our Parallel ϵ EP Algorithm to find the *coarsest equitable partition* of very large graphs for an $\epsilon = 0$. Finding the equitable partition of a graph forms an important intermediate stage in all the practical graph automorphism finding softwares [16, 5]. Another possible research direction is to explore algorithms for positional analysis of very large graphs using vertex-centric computation paradigms such as Pregel and GraphChi [14, 12].

6 Acknowledgements

We thank Inkit Padhi, Intern, RISE Lab, IIT Madras for helping prototype the code for the Fast ϵ EP algorithm. We would also like to thank Srikanth R. Madikeri, Research Scholar, DON Lab, IIT Madras for his valuable inputs during discussions on implementing the Parallel ϵ EP code.

References

- [1] S.P. BORGATTI AND M.G. EVERETT, *Notions of Position in Social Network Analysis*, Sociological Methodology, 22 (1992), pp. 1-35.
- [2] YINGYI BU, BILL HOWE, MAGDALENA BALAZINSKA, AND MICHAEL D ERNST, *HaLoop: Efficient Iterative*

- Data Processing on Large Clusters*, Proceedings of the VLDB Endowment, 3 (2010), pp. 285–296.
- [3] A. CARDON AND M. CROCHEMORE, *Partitioning a Graph in $O(|A|\log_2|V|)$* , Theoretical Computer Science, 19 (1982), pp. 85–98.
- [4] AARON CLAUSET, COSMA ROHILLA SHALIZI, AND MARK EJ NEWMAN, *Power-Law Distributions in Empirical Data*, SIAM review, 51 (2009), pp. 661–703.
- [5] PAUL T DARGA, KAREM A SAKALLAH, AND IGOR L MARKOV, *Faster Symmetry Discovery using Sparsity of Symmetries*, in Proceedings of the 45th Annual Design Automation Conference, ACM, 2008, pp. 149–154.
- [6] JEFFREY DEAN AND SANJAY GHEMAWAT, *MapReduce: Simplified Data Processing on Large Clusters*, Communications of the ACM, 51 (2008), pp. 107–113.
- [7] LUCILE DENGÈUD AND ALAIN GUÉNOCHE, *Comparison of distance indices between partitions*, in Data Science and Classification, Springer, 2006, pp. 21–28.
- [8] JALIYA EKANAYAKE, HUI LI, BINGJING ZHANG, THILINA GUNARATHNE, SEUNG-HEE BAE, JUDY QIU, AND GEOFFREY FOX, *Twister: A Runtime for Iterative MapReduce*, in Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing, ACM, 2010, pp. 810–818.
- [9] MARTIN G EVERETT, *Role Similarity and Complexity in Social Networks*, Social Networks, 7 (1985), pp. 353–359.
- [10] APACHE SOFTWARE FOUNDATION, *Apache Hadoop*. Available: <http://hadoop.apache.org/>. Accessed February 7, 2014.
- [11] KIRAN KATE AND BALARAMAN RAVINDRAN, *Epsilon Equitable Partition: A Positional Analysis method for Large Social Networks*, in Proceedings of 15th International Conference on Management of Data, 2009.
- [12] AAO KYROLA, GUY BLELLOCH, AND CARLOS GUESTRIN, *GraphChi: Large-scale Graph Computation on just a PC*, in Proceedings of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 2012, pp. 31–46.
- [13] F. LORRAIN AND H.C. WHITE, *Structural Equivalence of Individuals in Social Networks*, Journal of Mathematical Sociology, 1 (1971), pp. 49–80.
- [14] GRZEGORZ MALEWICZ, MATTHEW H AUSTERN, AART JC BIK, JAMES C DEHNERT, ILAN HORN, NATY LEISER, AND GRZEGORZ CZAJKOWSKI, *Pregel: A System for Large-scale Graph Processing*, in Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data, ACM, 2010, pp. 135–146.
- [15] B. D. MCKAY, *Practical Graph Isomorphism*, Congressus Numerantium, 30 (1981).
- [16] BRENDAN D. MCKAY, *nauty User’s Guide (Version 2.4)*. <http://cs.anu.edu.au/bdm/nauty/>, 2009.
- [17] TOMASZ MICHALAK, KV AADITHA, PL SZCZEPAN-SKI, BALARAMAN RAVINDRAN, AND NICHOLAS R JENNINGS, *Efficient Computation of the Shapley Value for Game-Theoretic Network Centrality*, Journal of AI Research, 46 (2013), pp. 607–650.
- [18] ALAN MISLOVE, HEMA SWETHA KOPPULA, KRISHNA P GUMMADI, PETER DRUSCHEL, AND BOBBY BHATTACHARJEE, *Growth of the Flickr Social Network*, in Proceedings of the first workshop on Online Social Networks, ACM, 2008, pp. 25–30.
- [19] PRASHANTH MUNDKUR, VILLE TUULOS, AND JARED FLATOW, *Disco: A Computing Platform for Large-Scale Data Analytics*, in Proceedings of the 10th ACM SIGPLAN workshop on Erlang, ACM, 2011, pp. 84–89.
- [20] ROBERT PAIGE AND ROBERT E TARJAN, *Three Partition Refinement Algorithms*, SIAM Journal on Computing, 16 (1987), pp. 973–989.
- [21] SIDDHARTH SURI AND SERGEI VASSILVITSKII, *Counting Triangles and the Curse of the Last Reducer*, in Proceedings of the 20th International Conference on World Wide Web, ACM, 2011, pp. 607–614.
- [22] OLE TANGE, *GNU Parallel - The Command-line Power Tool*, login: The USENIX Magazine, (2011), pp. 42–47.
- [23] ANDREW TRIDGELL AND PAUL MACKERRAS, “*The rsync algorithm*,” Australian National University, tech. report, TR-CS-96-05, 1996.
- [24] BIMAL VISWANATH, ALAN MISLOVE, MEEYOUNG CHA, AND KRISHNA P GUMMADI, *On the Evolution of User Interaction in Facebook*, in Proceedings of the 2nd ACM Workshop on Online Social Networks, ACM, 2009.
- [25] STANLEY WASSERMAN AND CAROLYN ANDERSON, *Stochastic a posteriori Blockmodels: Construction and Assessment*, Social Networks, 9 (1987), pp. 1–36.
- [26] STANLEY WASSERMAN AND KATHERINE FAUST, *Social Network Analysis: Methods and Applications*, Cambridge University Press, 1994.
- [27] D. R. WHITE AND K. REITZ, *Graph and Semigroup Homomorphisms on Semigroups of Relations*, Social Networks, 5 (1983), pp. 193–234.
- [28] MATEI ZAHARIA, MOSHARAF CHOWDHURY, TATHAGATA DAS, ANKUR DAVE, JUSTIN MA, MURPHY MCCAULEY, MICHAEL J FRANKLIN, SCOTT SHENKER, AND ION STOICA, *Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing*, in Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation, USENIX Association, 2012, pp. 2–2.
- [29] MATEI ZAHARIA, MOSHARAF CHOWDHURY, MICHAEL J FRANKLIN, SCOTT SHENKER, AND ION STOICA, *Spark: Cluster Computing with Working Sets*, in Proceedings of the 2nd USENIX Conference on Hot topics in Cloud Computing, 2010.

A Partition Similarity Score

A.1 Mathematical Preliminaries This section briefs out few mathematical preliminaries, which form the basis for our partition similarity score.

Given, graph $G \equiv \langle V, E \rangle$, V is the vertex set, E is the edge set and $\pi = \{C_1, C_2, \dots, C_K\}$ is a partition of V . We define the following for any two partitions of a

graph $G \equiv \langle V, E \rangle$:

- (i) Two *partitions* π_1 and π_2 are **equal**, *iff* they both partition the vertex set V of G exactly in the same way of each other.

Example,

$$\begin{aligned} \text{let, } \pi_1 &= \{\{v_1, v_2, v_3, v_4\}, \{v_5, v_6\}, \{v_7\}, \{v_8, v_9, v_{10}\}\} \\ \pi_2 &= \{\{v_6, v_5\}, \{v_3, v_2, v_4, v_1\}, \{v_9, v_8, v_{10}\}, \{v_7\}\}. \end{aligned}$$

$$\text{then, } \pi_1 = \pi_2$$

i.e. the order of cells in partition and the order of vertices inside a cell is not important.

- (ii) We define the **intersection** of two *partitions* π_1 and π_2 , as a partition containing the cells obtained from the *set intersection* operator applied *cell-wise* to member cells of π_{ϵ_1} and π_{ϵ_2} (discarding the *empty* sets).

Example,

$$\text{let, } \pi_1 = \{\{v_1, v_2, v_3\}, \{v_4, v_5\}, \{v_6, v_7, v_8\}\} \text{ and}$$

$$\pi_2 = \{\{v_1, v_2\}, \{v_3, v_4, v_5\}, \{v_6, v_7\}, \{v_8\}\}.$$

$$\text{then, } \pi_1 \cap \pi_2 = \{\{v_1, v_2\}, \{v_3\}, \{v_4, v_5\}, \{v_6, v_7\}, \{v_8\}\}$$

- (iii) Two *partitions* π_1 and π_2 are **dissimilar**, *iff* their intersection leads to a *discrete* partition. A discrete partition is a one with only *singleton* cells.

Example,

$$\text{let, } \pi_1 = \{\{v_1, v_2, v_3\}, \{v_4, v_5\}\} \text{ and}$$

$$\pi_2 = \{\{v_1, v_4\}, \{v_3, v_5\}, \{v_2\}\}.$$

$$\text{then, } \pi_1 \cap \pi_2 = \{\{v_1\}, \{v_2\}, \{v_3\}, \{v_4\}, \{v_5\}\}$$

Here, $\pi_1 \cap \pi_2$ gives a *discrete* partition.

A.2 Simplified Representation of the Partition Similarity Score Equation 4.4 can be represented in a simplified form as follows:

$$\begin{aligned} \text{(A.1) } \text{sim}(\pi_1, \pi_2) &= \frac{1}{2} \left[\left(\frac{N - |\pi_1 \cap \pi_2|}{N - |\pi_1|} \right) + \left(\frac{N - |\pi_1 \cap \pi_2|}{N - |\pi_2|} \right) \right] \\ &= \frac{1}{2} \left(1 - \frac{|\pi_1 \cap \pi_2|}{N} \right) \left[\left(\frac{1}{1 - \frac{|\pi_1|}{N}} \right) + \left(\frac{1}{1 - \frac{|\pi_2|}{N}} \right) \right] \\ &= \frac{1}{2} C(\pi_1 \cap \pi_2) \left[\frac{1}{C(\pi_1)} + \frac{1}{C(\pi_2)} \right] \\ &= \frac{C(\pi_1 \cap \pi_2)}{H(C(\pi_1), C(\pi_2))} \end{aligned}$$

Where,

$$C(\pi) = \left(1 - \frac{|\pi|}{N} \right),$$

$|\pi|$ = cardinality of π ,

N = cardinality of the *discrete* partition of π ,

$H(x, y)$ = *harmonic mean* of x and y .

The authors in [7] survey and compare several notions of distance indices between partitions on the same set, which are available in the literature.

A.3 MapReduce Algorithm to Compute the Partition Similarity Score The partition similarity score of Equation 4.4 requires the cardinality of the intersection set of the two partitions π_1 and π_2 .

Finding the intersection of two partitions as per the definition of *intersection* from (ii), Appendix A.1 is $O(n^2)$ operation, where n being the total number of vertices in the partition. Computing this for very large graphs becomes intractable. To counter this problem, we provide an algorithm based on the MapReduce paradigm [6] to compute the size of the *intersection set* of π_1 and π_2 (*i.e.*, $|\pi_1 \cap \pi_2|$). The algorithm is presented in Algorithm box 3. The algorithm initializes by *enumerating* the indices of π_2 for each cell index of partition π_1 . For each key from this tuple, the MAP operation checks if these two cells *intersect* or not. The MAP emits a *value* of 1 corresponding to a *constant key*. The REDUCE operations computes the *sum* of these individual 1s. This *sum* corresponds to $|\pi_1 \cap \pi_2|$, which is used to compute the partition similarity score from Equation 4.4.

A note on Algorithm 3: The INITIALIZE method of Algorithm 3 (*line* 3), primarily involves replicating/enumerating the cell indices of π_2 to all the cell indices of π_1 . Since *cross* operations are computationally very costly, the tractability of the Algorithm is inherently dependant on ability to generate the *cross product* set of the cell index tuples of the two input partitions.

Algorithm 3 MapReduce Partitions Intersection Set Cardinality

Input: Partitions π_1 and π_2 . Let $\pi_1 = \{cell_1:[v_1, v_2], cell_2:[v_3, v_4, v_5], \dots, cell_K:[v_{n-2}, v_{n-1}, v_n]\}$
and $\pi_2 = \{cell_1:[v_1, v_2, v_3], cell_2:[v_4, v_5], \dots, cell_L:[v_{n-1}, v_n]\}$
Output: Partitions intersection set cardinality $|\pi_1 \cap \pi_2|$

```
1 class MAPPER
2   tupleList enumCells=[]

3   method INITIALIZE()
4     for each cell_index i of  $\pi_1$  (i.e.  $1 \rightarrow K$ )
5       enumerate the cell_index j of  $\pi_2$  (i.e.  $1 \rightarrow L$ ) do
6         add tuple(i, j) them to enumCells

7   method MAP(id t, tuple (i, j))
8     intersect  $\leftarrow \{\pi_1(i) \cap \pi_2(j)\}$   $\triangleright$  Appendix A.1,
9     (ii)
10    if intersect  $\neq \phi$  then
11      EMIT(id intersect, 1)  $\triangleright$  If the two cells have
12      a overlap, emit value 1 corresponding to a constant
13      key "intersect"
```

```
1 class REDUCER
2   method REDUCE(id key, values)
3     sum = 0
4     for value in values
5       sum = sum + value
6     EMIT(key, sum)  $\triangleright |\pi_1 \cap \pi_2|$ 
```
