

Parallel Toolkit for Measuring the Quality of Network Community Structure

Mingming Chen

Department of Computer Science
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180
Email: chenm8@rpi.edu

Sisi Liu

Information Technology and Web Science
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180
Email: lius10@rpi.edu

Boleslaw K. Szymanski

Department of Computer Science
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180
Email: szymab@rpi.edu

Abstract—Many networks display community structure which identifies groups of nodes within which connections are denser than between them. Detecting and characterizing such community structure, which is known as community detection, is one of the fundamental issues in the study of network systems. It has received a considerable attention in the last years. Numerous techniques have been developed for both efficient and effective community detection. Among them, the most efficient algorithm is the label propagation algorithm whose computational complexity is $O(|E|)$. Although it is linear in the number of edges, the running time is still too long for very large networks, creating the need for parallel community detection. Also, computing community quality metrics for community structure is computationally expensive both with and without ground truth. However, to date we are not aware of any effort to introduce parallelism for this problem. In this paper, we provide a parallel toolkit¹ to calculate the values of such metrics. We evaluate the parallel algorithms on both distributed memory machine and shared memory machine. The experimental results show that they yield a significant performance gain over sequential execution in terms of total running time, speedup, and efficiency.

I. INTRODUCTION

Many networks, including Internet, citation networks, transportation networks, e-mail networks, and social and biochemical networks, display community structure which identifies groups of nodes within which connections are denser than between them [1]. Detecting and characterizing such community structure, which is known as community detection, is one of the fundamental issues in the study of network systems. Community detection has been shown to reveal latent yet meaningful structures in networks [2].

Thus, numerous techniques were developed for both efficient and effective community detection, including Modularity Optimization [3], [4], Clique Percolation [5], Local Expansion [6], [7], Fuzzy Clustering [8], [9], Link Partitioning [10], and Label Propagation [11], [12], [13]. Among them, the most efficient algorithm is the label propagation algorithm whose computational complexity is $O(|E|)$, where $|E|$ is the number of edges in the network. Although it is a linear algorithm, the running time is still too long for very large networks. The primary examples are online social networks that are increasingly popular, the largest being Facebook with more than 800 million daily active users². The WWW forms a

network of hyperlinked webpages in excess of 30 billion nodes. Therefore, parallelism was introduced into community detection to alleviate computational costs [14], [15], [16], [17], [18]. However, to date we are not aware of any effort that provides parallel computation for the community quality metrics with and without ground truth community structure [19], [20], [21], though it is computational expensive to do so. Hence, in this paper, we provide a parallel toolkit to calculate the values of these metrics. Although we are using parallel computing to speed up the processing, in most of the cases, algorithms are highly parallelizable, so the contributions of this paper focus on making the highly efficient social network analysis tools available to research community.

We implement the parallel algorithms with MPI (Message Passing Interface) and Pthreads (POSIX Threads). We perform runs on both distributed memory machine, such as Blue Gene/Q, and shared memory machine, like GANXIS. The network we adopt is LFR benchmark network [22]. The LFR benchmark network for testing the parallel programs to calculate the metrics with ground truth community structure has 100,000 nodes. We choose two sizes, ten million of nodes (10,000,000) and one hundred million of nodes (100,000,000), of the LFR benchmark network to test the parallel programs for computing the metrics without ground truth communities. The experimental results show that both parallel MPI algorithms and parallel Pthreads algorithms yield a significant performance gain over sequential execution. Moreover, we recommend using parallel MPI algorithms and parallel Pthreads algorithm respectively to calculate the metrics with and without ground truth communities on GANXIS (or shared memory machines) in terms of their speedup and efficiency.

II. COMMUNITY QUALITY METRICS

A. Metrics with Ground Truth Communities

The quality evaluation metrics with ground truth community structure we consider here can be divided into three categories: *Variation of Information (VI)* and *Normalized Mutual Information (NMI)* based on information theory; *F-measure* and *Normalized Van Dongen metric (NVD)* based on cluster matching; *Rand Index (RI)*, *Adjusted Rand Index (ARI)*, and *Jaccard Index (JI)* based on pair counting [19].

1) *Information Theory Based Metrics*: Given partitions C and C' , VI quantifies the “distance” between those two partitions, while NMI measures the similarity between them.

¹Please contact Mingming Chen via mileschen2008@gmail.com for the parallel toolkit if you are interested in it.

²Facebook company info: <http://newsroom.fb.com/company-info/>

VI is defined as

$$VI(C, C') = -\frac{1}{|V|} \sum_{c \in C, c' \in C'} |c \cap c'| \log \left(\frac{|c \cap c'|^2}{|c| |c'|} \right), \quad (1)$$

and NMI is given by

$$NMI(C, C') = \frac{-2 \sum_{c \in C, c' \in C'} \frac{|c \cap c'|}{|V|} \log \left(\frac{|c \cap c'| |V|}{|c| |c'|} \right)}{\sum_{c \in C} \frac{|c|}{|V|} \log \frac{|c|}{|V|} + \sum_{c' \in C'} \frac{|c'|}{|V|} \log \frac{|c'|}{|V|}}, \quad (2)$$

where $|V|$ is the number of nodes in the network, $|c|$ is the number of nodes in community c of C , and $|c \cap c'|$ is the number of nodes both in community c of C and in community c' of C' . The computational complexity to calculate VI and NMI is $O(|V||C'|)$, where $|C'|$ is the number of communities found by a community detection algorithm.

2) *Cluster Matching Based Metrics*: Measures based on cluster matching aim at finding the largest overlaps between pairs of communities of two partitions C and C' . *F-measure* measures the similarity between two partitions, while *NVD* quantifies the ‘‘distance’’ between partitions C and C' . *F-measure* is defined as

$$F\text{-measure}(C, C') = \frac{1}{|V|} \sum_{c \in C} |c| \max_{c' \in C'} \frac{2|c \cap c'|}{|c| + |c'|}. \quad (3)$$

NVD is given by

$$NVD(C, C') = 1 - \frac{1}{2|V|} \left(\sum_{c \in C} \max_{c' \in C'} |c \cap c'| + \sum_{c' \in C'} \max_{c \in C} |c' \cap c| \right). \quad (4)$$

The complexity to calculate *F-measure* and *NVD* is $O(|V|(|C| + |C'|))$, where $|C|$ is the number of communities in the ground truth community structure.

3) *Pair Counting Based Metrics*: Metrics based on pair counting count the number of pairs of nodes that are classified (in the same community or in different communities) in two partitions C and C' . Let a_{11} indicate the number of pairs of nodes that are in the same community in both partitions, a_{10} denote the number of pairs of nodes that are in the same community in C but in different communities in C' , a_{01} be the number of pairs of nodes which are in different communities in C but in the same community in C' , a_{00} be the number of pairs of nodes which are in different communities in both partitions. By definition, $A = a_{11} + a_{10} + a_{01} + a_{00} = \frac{|V|(|V|-1)}{2}$ is the total number of pairs of nodes in the network. Then, RI which is the ratio of the number of node pairs placed in the same way in both partitions to the total number of pairs is given by

$$RI(C, C') = \frac{a_{11} + a_{00}}{A}. \quad (5)$$

Denote $M = \frac{1}{A}(a_{11} + a_{10})(a_{11} + a_{01})$. Then, RI 's corresponding adjusted version, ARI , is expressed as

$$ARI(C, C') = \frac{a_{11} - M}{\frac{1}{2} [(a_{11} + a_{10}) + (a_{11} + a_{01})] - M}. \quad (6)$$

JI which is the ratio of the number of node pairs placed in the same community in both partitions to the number of node pairs that are placed in the same group in at least one partition is defined as

$$JI(C, C') = \frac{a_{11}}{a_{11} + a_{10} + a_{01}}. \quad (7)$$

Each of these three metrics quantifies the similarity between two partitions C and C' . The complexity to calculate RI , ARI , and JI is $O(|V|^2)$.

B. Metrics without Ground Truth Communities

1) *Newman's Modularity*: Modularity [1] measures the difference between the actual fraction of edges within the community and such fraction expected in a randomized graph with the same number of nodes and the same degree sequence. For the given community partition of a unweighted and undirected graph $G = (V, E)$ with $|E|$ edges, modularity (Q) is given by

$$Q = \sum_{c \in C} \left[\frac{|E_c^{in}|}{|E|} - \left(\frac{2|E_c^{in}| + |E_c^{out}|}{2|E|} \right)^2 \right], \quad (8)$$

where C is the set of all the communities, c is a specific community in C , $|E_c^{in}|$ is the number of edges between nodes within community c , and $|E_c^{out}|$ is the number of edges from the nodes in community c to the nodes outside c .

2) *Modularity Density*: *Modularity Density* (Q_{ds}) [20], [21] is proposed to solve the two opposite yet coexisting problems of modularity: in some cases, it tends to favor small communities over large ones while in others, large communities over small ones. The latter tendency is also known as the resolution limit problem [23]. For unweighted and undirected networks, Q_{ds} is defined as

$$Q_{ds} = \sum_{c_i \in C} \left[\frac{|E_{c_i}^{in}|}{|E|} d_{c_i} - \left(\frac{2|E_{c_i}^{in}| + |E_{c_i}^{out}|}{2|E|} d_{c_i} \right)^2 \right] - \sum_{\substack{c_j \in C \\ c_j \neq c_i}} \frac{|E_{c_i, c_j}|}{2|E|} d_{c_i, c_j}, \quad (9)$$

$$d_{c_i} = \frac{2|E_{c_i}^{in}|}{|c_i|(|c_i| - 1)},$$

$$d_{c_i, c_j} = \frac{|E_{c_i, c_j}|}{|c_i| |c_j|}.$$

In the above, d_{c_i} is the internal density of community c_i , $|E_{c_i, c_j}|$ is the number of edges from c_i to c_j , and d_{c_i, c_j} is the pair-wise density between communities c_i and c_j .

3) *Six Other Community Quality Measures*: We also consider six other metrics without ground truth community structure, including the number of *Intra-edges*, *Intra-density*, *Contraction*, the number of *Inter-edges*, *Expansion*, and *Conductance* [20], [21], which characterize how community-like is the connectivity structure of a given set of nodes. All of them rely on the intuition that communities are sets of nodes with many edges inside and few edges outside.

The number of Intra-edges: $|E_c^{in}|$; it is the total number of edges in community c .

Intra-density: d_{c_i} in Equation (9).

Contraction: $2|E_c^{in}|/|c|$; it measures the average number of edges per node inside the community c .

The number of Inter-edges: $|E_c^{out}|$; it is the total number of edges on the boundary of c .

Expansion: $|E_c^{out}|/|c|$; it measures the average number of edges (per node) that point outside the community c .

Conductance: $\frac{|E_c^{out}|}{2|E_c^{in}| + |E_c^{out}|}$; it measures the fraction of the total number of edges that point outside the community.

III. PARALLEL ALGORITHM DESIGN

In this section, we present the parallel algorithms, MPI and Pthreads versions, to calculate the quality metrics introduced in Section II. It can be seen from Section II that VI and NMI can be calculated together, F -measure and NVD can be computed together, RI , ARI , and JI can be calculated together, and the metrics without ground truth communities can be computed together. Thus, we will have four parallel algorithms based on MPI and four parallel algorithms based on Pthreads for these metrics. We denote the ground truth community structure as C and the community structure detected with a community detection algorithm as C' .

A. Parallel Algorithms Based on MPI

In the parallel algorithms based on MPI, the problem to calculate metrics is partitioned with the unit of community. For the algorithms to calculate the metrics with ground truth community structure, each processor extracts the ground truth communities and the detected communities when $(comId \bmod numProcs == procId)$ to achieve rough load balance. $comId$ is the id of a community, $procId$ is the id of a processor, and $numProcs$ is the total number of processors used. Hence, each processor will have $\lfloor |C|/numProcs \rfloor$ or $\lfloor |C|/numProcs \rfloor + 1$ ground truth communities and $\lfloor |C'|/numProcs \rfloor$ or $\lfloor |C'|/numProcs \rfloor + 1$ discovered communities. For the algorithms to compute the metrics without ground truth communities, each processor extracts the discovered communities using the same approach. Also, each processor gets its local network which contains the nodes in its own communities and the neighboring nodes of these nodes.

Algorithm 1 MPI_information_theory_metric(C_p, C'_p)

```

1: // Processor id is denoted as  $p$ . The number of processors
   used is denoted as  $numProcs$ . The local values of  $VI$ 
   and  $NMI$  is denoted as  $rankVI$  and  $rankNMI$ .
2: Calculate  $rankVI$  and  $rankNMI$  based on Equation (1)
   and Equation (2) using its own  $C_p$  and its own  $C'_p$ ;
3: // Circulate  $C'_p$  in a ring.
4:  $recvSrc = (p + numProcs - 1) \bmod numProcs$ ;
5:  $sendDst = (p + 1) \bmod numProcs$ ;
6:  $receivedMsgNum = 0$ ;
7: while  $receivedMsgNum < (numProcs - 1)$  do
8:    $sendBuf[] \leftarrow C'_p$ ;
9:   Send  $sendBuf[]$  to  $sendDst$ ;
10:  Receive  $C'_p$  from  $recvSrc$  with  $recvBuf[]$ ;
11:  Calculate  $rankVI$  and  $rankNMI$  based on Equations
     (1) and (2) using its own  $C_p$  and received  $C'_p$ ;
12:   $++receivedMsgNum$ ;
13: end while
14: Get  $VI$  and  $NMI$  by summing the  $rankVI$  and
      $rankNMI$  of all  $numProcs$  processors;
15: Return  $VI$  and  $NMI$ ;

```

We first show the parallel MPI algorithm to calculate the information theory based metrics, VI and NMI . Supposed there are N processors (or MPI ranks), each processor reads its own set of ground truth communities C_p and its own set of discovered communities C'_p . It can be learnt from the definitions of VI and NMI given by Equation (1) and Equation (2) respectively that each ground truth community

should traverse all the discovered communities in order to get the values of them. Thus, in the algorithm, we circulate C'_p to each processor in a “ring”. That is, processor 0 sends its C'_p to processor 1, and processor 1 to processor 2, and processor $N - 1$ would send C'_p to processor 0. This “shifting” of C'_p will occur $N - 1$ times for N processors. For its own C'_p and each received C'_p , the processor will calculate its local values of VI and NMI with its own C_p . Finally, the values of VI and NMI are the sum of the local values of all N processors. Algorithm 1 shows our parallel algorithm for computing VI and NMI based on MPI. It takes C_p and C'_p as parameters.

Algorithm 2 MPI_cluster_matching_metric(C_p, C'_p)

```

1: // Use  $maxNormedComs[]$  and  $maxTComs[]$  to record
   the max item for each ground truth community shown in
   Equations (3) and (4), respectively.
2: Get  $maxNormedComs[]$  and  $maxTComs[]$  for each
   community in  $C_p$  with its own  $C'_p$ ;
3:  $recvSrc = (p + numProcs - 1) \bmod numProcs$ ;
4:  $sendDst = (p + 1) \bmod numProcs$ ;
5: // Circulate  $C'_p$  in a ring.
6:  $receivedMsgNum = 0$ ;
7: while  $receivedMsgNum < (numProcs - 1)$  do
8:    $sendBuf[] \leftarrow C'_p$ ;
9:   Send  $sendBuf[]$  to  $sendDst$ ;
10:  Receive  $C'_p$  from  $recvSrc$  with  $recvBuf[]$ ;
11:  Update  $maxNormedComs[]$  and  $maxTComs[]$  for
     each community in  $C_p$  with received  $C'_p$ ;
12:   $++receivedMsgNum$ ;
13: end while
14: // Use  $maxDComs[]$  to record the max item for each
   detected community shown in Equation (4).
15: Get  $maxDComs[]$  for each community in  $C'_p$  with its own
   set of ground truth communities  $C_p$ ;
16: // Circulate  $C_p$  in a ring.
17:  $receivedMsgNum = 0$ ;
18: while  $receivedMsgNum < (numProcs - 1)$  do
19:    $sendBuf[] \leftarrow C_p$ ;
20:   Send  $sendBuf[]$  to  $sendDst$ ;
21:   Receive  $C_p$  from  $recvSrc$  with  $recvBuf[]$ ;
22:   Update  $maxDComs[]$  for each community in  $C'_p$  with
     received  $C_p$ ;
23:    $++receivedMsgNum$ ;
24: end while
25: Calculate  $rankFMeasure$  and also  $rankNVD$  with
      $maxNormedComs[]$ ,  $maxTComs[]$ , and  $maxDComs[]$ 
     based on Equations (3) and (4);
26: Get  $F$ -measure and  $NVD$  by summing  $rankFMeasure$ 
     and  $rankNVD$  of all  $numProcs$  processors;
27: Return  $F$ -measure and  $NVD$ ;

```

We then present the parallel algorithm to calculate the cluster matching based metrics, F -measure and NVD . From the definitions of F -measure and NVD shown in Equation (3) and Equation (4) respectively, we could learn that in order to calculate F -measure and NVD , we need to determine for each ground truth community the discovered community that has the largest number of common nodes with it. In addition, to calculate NVD , we further need to locate for each discovered community the ground truth community that has the largest number of common nodes with it. Hence, in the algorithm,

both C'_p and C_p are circulated to each processor in a “ring”. This “shifting” of C'_p and C_p will both occur $N-1$ times for N processors. For its own C'_p and each received C'_p , the processor will calculate its local values of F -measure and NVD with its own C_p . Moreover, for its own C_p and each received C_p , the processor will compute its local values of NVD with its own C'_p . Finally, the values of F -measure and NVD are the sum of the local values of all N processors. Algorithm 2 shows our parallel MPI algorithm for computing F -measure and NVD . It takes C_p and C'_p as parameters.

Algorithm 3 MPI_pair_counting_metric($C_p(map)$, $C'_p(map)$)

```

1: Count rankA11, rankA10, rankA01, and rankA00 using
   its own  $C_p(map)$  and its own  $C'_p(map)$ ;
2: // Circulate  $C'_p(map)$  in a ring.
3:  $recvSrc = (p + numProcs - 1) \bmod numProcs$ ;
4:  $sendDst = (p + 1) \bmod numProcs$ ;
5:  $receivedMsgNum = 0$ ;
6: while  $receivedMsgNum < (numProcs - 1)$  do
7:    $sendBuf[] \leftarrow C'_p(map)$ ;
8:   Send  $sendBuf[]$  to  $sendDst$ ;
9:   Receive  $C'_p(map)$  from  $recvSrc$  with  $recvBuf[]$ ;
10:  Count rankA01 and rankA00 using its own  $C_p(map)$ 
    and received  $C'_p(map)$ ;
11:   $++receivedMsgNum$ ;
12: end while
13: Calculate  $rankRI$ ,  $rankARI$ , and  $rankJI$  with
     $rankA11$ ,  $rankA10$ ,  $rankA01$ , and  $rankA00$ ;
14: Get  $RI$ ,  $ARI$ , and  $JI$  by summing the  $rankRI$ ,
     $rankARI$ , and  $rankJI$  of all  $numProcs$  processors;
15: Return  $RI$ ,  $ARI$ , and  $JI$ ;

```

We now demonstrate how to calculate the pair counting based metrics, RI , ARI , and JI , in parallel with MPI. To calculate RI , ARI , and JI , each node in the network needs to traverse all the other nodes so as to get a_{11} , a_{10} , a_{01} , and a_{00} . Therefore, each processor reads its own ground truth communities and saves as a map with key being the node id and value being the id of the ground truth community to which this node belongs. We denote the map of nodes with their communities from ground truth community structure as $C_p(map)$. This processor also reads the community information for the nodes in $C_p(map)$ from discovered community structure and saves as a map. We denote the map of nodes with their communities from discovered community structure as $C'_p(map)$. $C_p(map)$ and $C'_p(map)$ have the same subset of nodes but with their community information from ground truth community structure and detected community structure, respectively. In our implementation, we use hash indexed map in order to search the community that a node belongs to quickly. Since each node needs to traverse all the other nodes, thus in the algorithm $C'_p(map)$ is circulated to each processor in a “ring”. This “shifting” of $C'_p(map)$ will also occur $N-1$ times for N processors. For each processor, each node in $C_p(map)$ first traverses the other nodes in its own $C'_p(map)$ to count its local values of a_{11} , a_{10} , a_{01} , and a_{00} . Then, this node will traverse the nodes in received $C'_p(map)$ to count only a_{01} and a_{00} because this node is in a different ground truth community with the nodes in received $C'_p(map)$. With a_{11} , a_{10} , a_{01} , and a_{00} , the processor could get the local values of RI , ARI , and JI based on Equations (5), (6), and (7).

Finally, the values of RI , ARI , and JI are the sum of the local values of all N processors. Algorithm 3 shows our parallel MPI algorithm for computing RI , ARI , and JI . It takes $C_p(map)$ and $C'_p(map)$ as parameters.

Finally, we illustrate how to calculate the community quality metrics without ground truth community structure, such as modularity and *Modularity Density*, in parallel with MPI. From Section II-B, we could observe that in order to calculate the contribution of a community to these metrics, we only need to obtain the number of edges inside it, the number of edges on the boundary of it, the numbers of edges between it and its neighboring communities, and the sizes of it and its neighboring communities. There is no dependency between processors. Hence, there is no need to transfer communities or to transfer any message between processors. In the algorithm, each processor reads its own set of discovered communities C'_p and its local network, and then calculate its own part for these metrics. At last, the values of these metrics are the sum of the local values of all N processors. We will not show the outline of this algorithm here because of its simplicity.

Algorithm 4 Pthreads_pair_counting_metric($nodes$)

```

1: // Thread id is denoted as  $threadId$ . The number of
   threads used is denoted as  $numThreads$ .
2: // The number of nodes in the network.
3:  $numNodes = nodes.size()$ ;
4: for  $i = 0$  to  $numNodes$  do
5:   if  $i \bmod numThreads == threadId$  then
6:      $iNode = nodes[i]$ ;
7:     // Traverse all the other nodes for  $iNode$ .
8:     for  $j = 0$  to  $numNodes$  do
9:        $jNode = nodes[j]$ ;
10:      if  $iNode \neq jNode$  then
11:        Count  $a_{11}$ ,  $a_{10}$ ,  $a_{01}$ , and  $a_{00}$  based on the
          community information of  $iNode$  and  $jNode$ 
          from ground truth community structure and dis-
          covered community structure;
12:      end if
13:    end for
14:  end if
15: end for
16: Calculate  $RI$ ,  $ARI$ , and  $JI$  with  $a_{11}$ ,  $a_{10}$ ,  $a_{01}$ , and  $a_{00}$ 
    based on Equations (5), (6), and (7);
17: Return  $RI$ ,  $ARI$ , and  $JI$ ;

```

B. Parallel Algorithms Based on Pthreads

The parallel Pthreads algorithms for all the metrics, except the ones based on pair counting, assign subsets of ground truth communities and discovered communities, and also local network to each thread using the same approach adopted in the parallel MPI algorithms introduced in Section III-A. The difference between the parallel Pthreads algorithms and the parallel MPI algorithms is that the ground truth communities, the detected communities, and the network are globally accessible in Pthreads, while they are locally stored in MPI. The cores of the algorithms to calculate these metrics do not change compared with the parallel MPI algorithms, so we will not present their outlines here.

In the parallel Pthreads algorithm to calculate the pair counting based metrics, the problem is partitioned with the

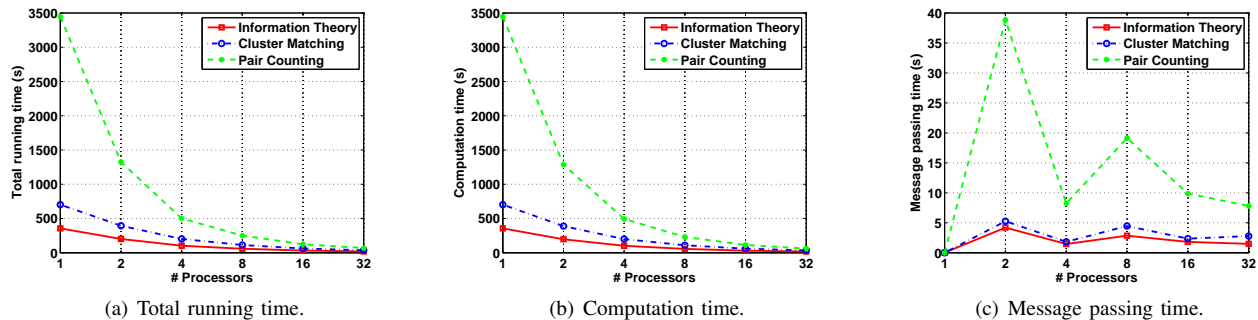


Fig. 1. The total running time, computation time, and message passing time of the three parallel MPI algorithms for computing the community quality metrics with ground truth community structure on GANXIS.

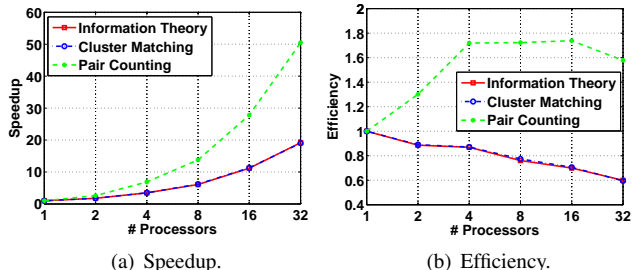


Fig. 2. The speedup and efficiency of the three parallel MPI algorithms for computing the metrics with ground truth community structure on GANXIS.

unit of node instead of the unit of community. Then, each node traverses all the other nodes in the network to count a_{11} , a_{10} , a_{01} , and a_{00} . The values of RI , ARI , and JI can be then calculated based on Equations (5), (6), and (7). The outline of this parallel Pthreads algorithm is shown in Algorithm 4.

IV. EVALUATION AND ANALYSIS

In this section, we first introduce the parallel architectures on which we perform runs for our parallel algorithms. Then, we introduce the measures that are used to evaluate the performance of these algorithms. We also give an introduction to LFR benchmark network [22] for which we calculate the metrics. Finally, we show the performance results of the parallel MPI and Pthreads algorithms.

A. Parallel Computing Architectures

We perform runs on both distributed memory machine, such as Blue Gene/Q, and shared memory machine, like GANXIS. We vary the number of processors used in GANXIS from 1 to 32 and the number of computing nodes (16 cores for each node) used in Blue Gene/Q from 1 to 256.

1) *GANXIS*: GANXIS a hyper threaded Linux system operating on a Silicon Mechanics Rackform nServ A422.v3 machine (GANXIS.nest.rpi.edu). Processing power was provided by 64 cores organized as four AMD OpteronTM 6272 (2.1 GHz, 16-core, G34, 16 MB L3 Cache) central processing units operating over a shared 512 GB of Random Access Memory (RAM) (32 x 16 GB DDR3-1600 ECC Registered 2R DIMMs) running at 1600 MT/s Max.

2) *Blue Gene/Q*: The Blue Gene/Q system that we used is stationed at The Computational Center for Nanotechnology Innovations facility at RPI, Troy, NY.

B. Performance Measures

We calculate *speedup* using Equation (10)

$$Speedup = T_1/T_p \quad (10)$$

where T_1 is the running time of the sequential program and T_p is the running time of the parallel program when p processors is adopted. We also compute *efficiency* according to Equation (11)

$$Efficiency = Speedup/p \quad (11)$$

where *Speedup* is the actual speedup calculated according to Equation (10) and p is the number of processors.

Notice that for our experimental results on Blue Gene/Q, p is denoted as the number of computing nodes adopted, T_1 is the running time of the parallel program when only 1 node adopted, and T_p is the running time of the parallel program when p nodes adopted.

C. LFR Benchmark Network

We run our parallel MPI and Pthreads programs to calculate the metrics for LFR benchmark networks [22] which have known ground truth community structure. The average node degree of the LFR benchmark networks is set to be 15 and the maximum node degree is set to be 50. The exponent γ for the degree sequence is 2. The exponent β for the community size distribution is 1. The mixing parameter μ is equal to 0.3.

The LFR benchmark network for testing the parallel programs for calculating the metrics with ground truth community structure has 100,000 nodes. The ground truth community structure is given when generating the network. The discovered community structure is obtained by using a community detection algorithm called Speaker listener Label Propagation Algorithm (SLPA) [12] with threshold parameter $r = 0.5$. SLPA gets disjoint communities when $r = 0.5$.

We choose two sizes, ten million of nodes (10,000,000) and one hundred million of nodes (100,000,000), of LFR network to test the parallel programs to compute the metrics without ground truth communities. We calculate the values of these metrics for the ground truth community structure instead of for the discovered community structure since it takes too long to get the detected communities with SLPA.

D. Experimental Results

In this part, we will report the performance results of the parallel MPI and Pthreads algorithms for community quality metrics both with and without ground truth community structure. For the running time, we do not take into account the I/O time. That is, the time of the program to read the ground truth communities, the discovered communities, or the network. Moreover, since GANXIS has 64 processors,

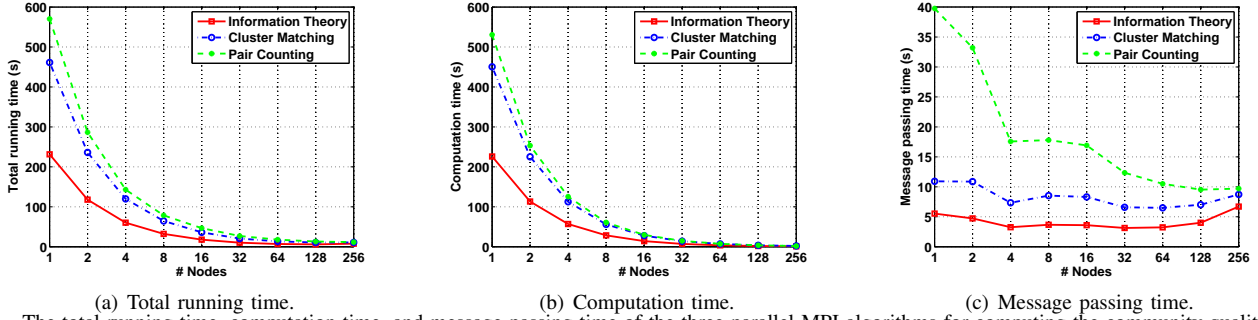


Fig. 3. The total running time, computation time, and message passing time of the three parallel MPI algorithms for computing the community quality metrics with ground truth community structure on Blue Gene/Q.

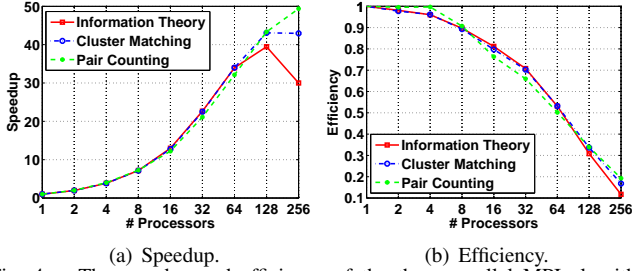


Fig. 4. The speedup and efficiency of the three parallel MPI algorithms for computing the community quality metrics with ground truth community structure on Blue Gene/Q.

every thread in our parallel Pthreads algorithms executes on its dedicated processor. Therefore, threads do not compete for central processing unit (CPU) processors. They execute in parallel, and we can completely ignore thread scheduling issues in our considerations. Because of this we use terms ‘thread’ and ‘processor’ interchangeably when describing the results of the parallel Pthreads algorithms on GANXIS.

1) *Performance of Parallel Algorithms for Metrics with Ground Truth Community Structure:* Figures 1(a), 1(b), and 1(c) respectively show the total running time, computation time, and message passing time of the three parallel MPI algorithms, Algorithm 1, Algorithm 2, and Algorithm 3, to compute the community quality metrics with ground truth communities on GANXIS. Figure 1(a) indicates that the total running time decreases as the number of processors increases. Figure 1(b) implies that the computation time decreases as the growth of the number of processors. However, it is shown in Figure 1(c) that the message passing time goes as a saw shape when the number of processors grows. The saw behavior is the result of GANXIS architecture in which quads of cores uses one shared global memory module. Increasing the number of processors from 1 to 2 introduces the message passing for the first time, adding two more processors helps as the message are placed in the same share memory module for all four cores. In case of 8 processors, some messages start to be moving between two memory modules, which decreases the advantage of having more processors and so on. In addition, Figures 2(a) and 2(b) present the corresponding speedup and efficiency. From Figure 2(a), we can observe that the speedup of all three algorithms grows as the number of processors increases. Figure 2(b) indicates that the efficiency of Algorithm 3 first increases from 1 to 16 processors and then decreases from 16 to 32 processors, while the efficiency of the other two algorithms always decreases. It is worth noting that the speedup and efficiency of Algorithm 1 and Algorithm 2 are

almost the same with each other, but both are much smaller than those of Algorithm 3 that is to calculate the pair counting based metrics. The efficiency of Algorithm 3 is even larger than 1, achieving a super-linear speedup. The super-linear speedup is the result of increasing larger cache available on many processors. As the number of processors increases, the volume of data processed on each processor decreases but cache is the same size. Thus, the number of cache misses decreases on each processor, speeding up the execution beyond linear speed up.

Figures 3(a), 3(b), and 3(c) respectively present the total running time, computation time, and message passing time of the three parallel MPI algorithms, Algorithm 1, Algorithm 2, and Algorithm 3, for computing the metrics with ground truth communities on Blue Gene/Q. Note that the x-axis is the number of nodes and each node has 16 processors. Thus, the number of processors is the number of nodes times 16. Figure 3(a) demonstrates that the total running time decreases as the number of nodes grows. Figure 3(b) implies that the computation time decreases as the growth of the number of nodes. However, there is no obvious trend of the message passing time. Moreover, Figures 4(a) and 4(b) show the corresponding speedup and efficiency. It can be observed from Figure 4(a) that the speedup of the three algorithms grows when the number of nodes increases, except for Algorithm 1 and Algorithm 2 when there are 256 nodes, the reason of which is that the message passing time instead of the computation time is the dominant part of the total running time at this case. Figure 4(b) implies that the efficiency of all three algorithms decreases as the growth of the number of nodes.

Figures 5(a), 5(b), and 5(c) respectively show the total running time, speedup, and efficiency of the three parallel Pthreads algorithms for computing the metrics with ground truth community structure on GANXIS. Figure 5(a) implies that the total running time decreases when the number of processors increases. Figure 5(b) indicates that the speedup increases as the growth of the number of processors. However, we could learn from Figure 5(c) that the efficiency first grows from 1 to 2 processors and then decreases from 2 to 32 processors. Comparing the total running time in Figure 5(a) and in Figure 1(a), we could see that the total running time of the three parallel Pthreads algorithms is generally larger than that of the three parallel MPI algorithms. Therefore, we recommend using the parallel MPI algorithms instead of the parallel Pthreads algorithms to calculate the metrics with ground truth community structure on GANXIS (or shared memory machines). Also, it is interesting that the speedup and efficiency of the parallel Pthreads algorithms to calculate the information theory and cluster matching based metrics shown

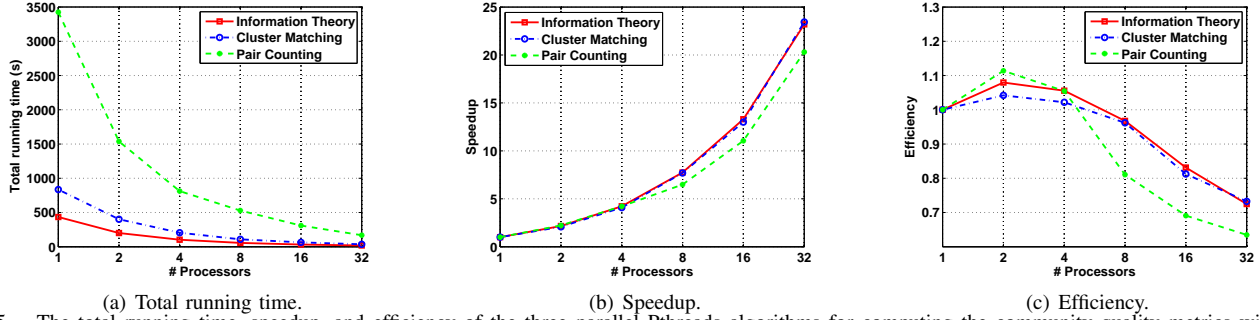


Fig. 5. The total running time, speedup, and efficiency of the three parallel Pthreads algorithms for computing the community quality metrics with ground truth community structure on GANXIS.

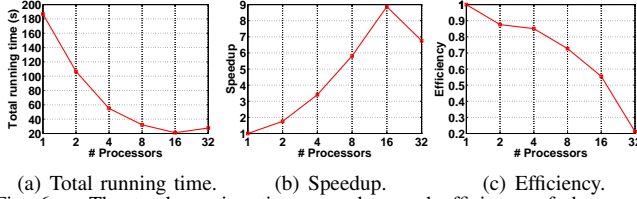


Fig. 6. The total running time, speedup, and efficiency of the parallel MPI algorithm for computing the community quality metrics without ground truth community structure on GANXIS (The number of nodes of the LFR benchmark network is 10,000,000.).

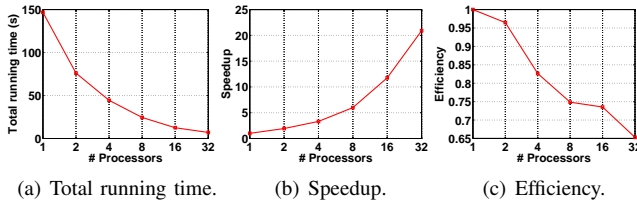


Fig. 7. The total running time, speedup, and efficiency of the parallel Pthreads algorithm for computing the community quality metrics without ground truth community structure on GANXIS (The number of nodes of the LFR benchmark network is 10,000,000.).

in Figures 5(b) and 5(c) are larger than those of the corresponding parallel MPI algorithms shown in Figures 2(a) and 2(b). However, the speedup and efficiency of the parallel Pthreads algorithm to calculate the pair counting based metrics shown in Figures 5(b) and 5(c) are much smaller than those of the corresponding parallel MPI algorithm shown in Figures 2(a) and 2(b). This phenomenon leads to an interesting result that the speedup and efficiency of the parallel Pthreads algorithm to calculate the pair counting based metrics are smaller than those of the other two parallel Pthreads algorithms, while the speedup and efficiency of the parallel MPI algorithm for calculating the pair counting based metrics are much larger than those of the other two parallel MPI algorithms.

2) *Performance of Parallel Algorithms for Metrics without Ground Truth Community Structure:* Figures 6(a), 6(b), and 6(c) respectively show the total running time, speedup, and efficiency of the parallel MPI algorithm for calculating the metrics without ground truth community structure on GANXIS with the size of the LFR benchmark network being 10,000,000. Figure 6(a) demonstrates that the total running time first decreases from 1 to 16 processors and then increases a little from 16 to 32 processors. Figure 6(b) indicates that the speedup first grows from 1 to 16 processors and then decreases from 16 to 32 processors. Thus, we can learn that there is a performance degradation when there are 32 processors. This performance penalty is again caused by the memory banks organization of

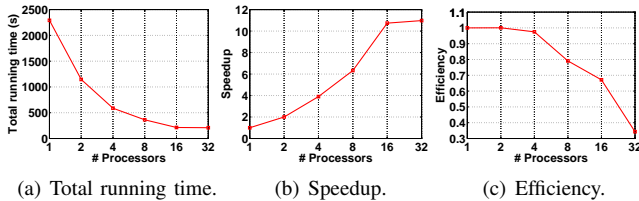
GANXIS machine. Also, it can be seen from Figure 6(c) that the efficiency always decreases.

Figures 7(a), 7(b), and 7(c) present the corresponding total running time, speedup, and efficiency of the parallel Pthreads algorithm. We could observe that the total running time decreases, the speedup grows, and the efficiency decreases as the number of processors increases. Also, the total running time, the speedup, and the efficiency of the parallel Pthreads algorithm are respectively much smaller, larger, and higher than those of the parallel MPI algorithm, compared to the results between Figure 6 and Figure 7.

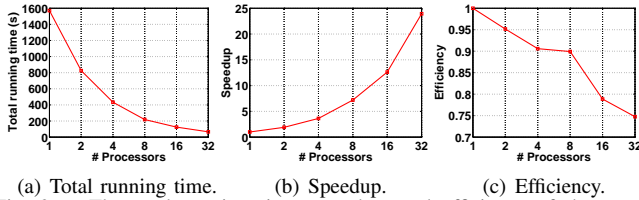
Figures 8(a), 8(b), and 8(c) respectively show the total running time, speedup, and efficiency of the parallel MPI algorithm for calculating the metrics without ground truth community structure on GANXIS with the size of the LFR benchmark network being 100,000,000. These three subfigures implies that the total running time decreases, the speedup increases, and the efficiency decreases as the growth of the number of processors. Similarly to the results shown in Figure 6, Figure 8 demonstrates that there is also a performance degradation when there are 32 processors.

Figures 9(a), 9(b), and 9(c) present the corresponding total running time, speedup, and efficiency of the parallel Pthreads algorithm. It shows that the total running time decreases, the speedup grows, and the efficiency decreases when the number of processors increases. Comparing the results between Figure 8 and Figure 9, we could observe that the total running time, the speedup, and the efficiency of the parallel Pthreads algorithm are respectively much smaller, larger, and higher than those of the parallel MPI algorithm.

When the size of the LFR benchmark network is 10,000,000, the smallest running time the parallel MPI algorithm achieves is 21.06 seconds at 16 processors, while the smallest running time the parallel Pthreads algorithm achieves is 7.02 seconds at 32 processors. The largest speedup the parallel MPI algorithm gets is 8.87 at 16 processors, while the largest speedup the parallel Pthread algorithm gets is 20.9 at 32 processors. When the size of the LFR benchmark network is 100,000,000, the smallest running time the parallel MPI algorithm can achieve is 208.82 seconds at 32 processors, while the smallest running time the parallel Pthreads algorithm can achieve is 65.66 seconds at 32 processors. The largest speedup the parallel MPI algorithm can get is 10.98 at 32 processors, while the largest speedup the parallel Pthread algorithm can get is 23.92 at 32 processors. Thus, we recommend using the parallel Pthreads algorithm instead of the parallel MPI algorithm to calculate the metrics without ground



(a) Total running time. (b) Speedup. (c) Efficiency.
 Fig. 8. The total running time, speedup, and efficiency of the parallel MPI algorithm for computing the community quality metrics without ground truth community structure on GANXIS (The number of nodes of the LFR benchmark network is 100,000,000.).



(a) Total running time. (b) Speedup. (c) Efficiency.
 Fig. 9. The total running time, speedup, and efficiency of the parallel Pthreads algorithm for computing the community quality metrics without ground truth community structure on GANXIS (The number of nodes of the LFR benchmark network is 100,000,000.).

truth communities on GANXIS (or shared memory machines). There is another point we could get is that the speedup and the efficiency of both parallel MPI algorithm and parallel Pthreads algorithm grow as the size of the network increases.

V. CONCLUSION

In this paper, we provide a parallel toolkit, implemented with MPI and Pthreads, to calculate the community quality metrics with and without ground truth community structure. We evaluate their performance on both distributed memory machine, such as Blue Gene/Q, and shared memory machine, for instance GANXIS. We conduct experiments on LFR benchmark networks with the number of nodes being 100,000; 10,000,000; and 100,000,000. The experimental results indicate that both the parallel MPI programs and the parallel Pthreads programs yield a significant performance gain over sequential execution. In addition, we discover that the parallel MPI algorithms perform better than the parallel Pthreads algorithms in terms of total running time, speedup, and efficiency on calculating the metrics with ground truth community structure, while the situation reverses on computing the metrics without ground truth community structure. Therefore, we recommend using the parallel MPI algorithms and the parallel Pthreads algorithm respectively to calculate the metrics with and without ground truth community structure.

ACKNOWLEDGMENT

This work was supported in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 and by the the Office of Naval Research Grant No. N00014-09-1-0607. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing the official policies either expressed or implied of the Army Research Laboratory or the U.S. Government.

REFERENCES

[1] M. E. J. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Phys. Rev. E*, vol. 69, p. 026113, Feb 2004.
 [2] S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, pp. 75–174, 2010.

[3] M. E. J. Newman, "Fast algorithm for detecting community structure in networks," *Phys. Rev. E*, vol. 69, p. 066133, Jun 2004.
 [4] M. E. J. Newman, "Modularity and community structure in networks," *Proceedings of the National Academy of Sciences*, vol. 103, no. 23, pp. 8577–8582, 2006.
 [5] G. Palla, I. Derényi, I. Farkas, and T. Vicsek, "Uncovering the overlapping community structure of complex networks in nature and society," *Nature*, no. 7043, pp. 814–818, 2005.
 [6] J. Baumes, M. Goldberg, and M. Magdon-ismail, "Efficient identification of overlapping communities," in *Proceedings of IEEE International Conference on Intelligence and Security Informatics*, 2005, pp. 27–36.
 [7] A. Lancichinetti, S. Fortunato, and J. Kertész, "Detecting the overlapping and hierarchical community structure in complex networks," *New Journal of Physics*, vol. 11, no. 3, p. 033015, 2009.
 [8] S. Zhang, R.-S. Wang, and X.-S. Zhang, "Identification of overlapping community structure in complex networks using fuzzy-means clustering," *Physica A: Statistical Mechanics and its Applications*, vol. 374, no. 1, pp. 483–490, 2007.
 [9] I. Psorakis, S. Roberts, M. Ebden, and B. Sheldon, "Overlapping community detection using bayesian non-negative matrix factorization," *Phys. Rev. E*, vol. 83, p. 066114, Jun 2011.
 [10] Y.-Y. Ahn, J. P. Bagrow, and S. Lehmann, "Link communities reveal multiscale complexity in networks," *Nature*, vol. 466, no. 7307, pp. 761–764, 2010.
 [11] U. N. Raghavan, R. Albert, and S. Kumara, "Near linear time algorithm to detect community structures in large-scale networks," *Phys. Rev. E*, vol. 76, p. 036106, Sep 2007.
 [12] J. Xie and B. K. Szymanski, "Towards linear time overlapping community detection in social networks," in *Proceedings of the 16th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2012, pp. 25–36.
 [13] J. Xie, M. Chen, and B. K. Szymanski, "LabelrankT: Incremental community detection in dynamic networks via label propagation," in *Proceedings of ACM SIGMOD Workshop on Dynamic Networks Management and Mining (DyNetMM)*, New York, USA, 2013.
 [14] Y. Zhang, J. Wang, Y. Wang, and L. Zhou, "Parallel community detection on large networks with propinquity dynamics," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA, 2009, pp. 997–1006.
 [15] E. J. Riedy, H. Meyerhenke, D. Ediger, and D. A. Bader, "Parallel community detection for massive graphs," in *Proceedings of the 9th International Conference on Parallel Processing and Applied Mathematics*, Berlin, Heidelberg, 2012, pp. 286–296.
 [16] E. Gregori, L. Lenzi, and S. Mainardi, "Parallel k-clique community detection on large-scale networks," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 24, no. 8, pp. 1651–1660, Aug 2013.
 [17] K. Kuzmin, S. Y. Shah, and B. K. Szymanski, "Parallel overlapping community detection with SLPA," in *Proceedings of the International Conference on Social Computing*, Washington, DC, USA, September 2013, pp. 204–212.
 [18] C. Staudt and H. Meyerhenke, "Engineering high-performance community detection heuristics for massive graphs," in *Proceedings of the International Conference on Parallel Processing*, 2013, pp. 180–189.
 [19] M. Chen, K. Kuzmin, and B. K. Szymanski, "Community detection via maximization of modularity and its variants," *Computational Social Systems, IEEE Transactions on*, vol. 1, no. 1, pp. 46–65, March 2014.
 [20] M. Chen, T. Nguyen, and B. K. Szymanski, "On measuring the quality of a network community structure," in *Proceedings of ASE/IEEE International Conference on Social Computing*, Washington, DC, USA, September 2013, pp. 122–127.
 [21] M. Chen, T. Nguyen, and B. K. Szymanski, "A new metric for quality of network community structure," *ASE Human Journal*, vol. 2, no. 4, pp. 226–240, 2013.
 [22] A. Lancichinetti, S. Fortunato, and F. Radicchi, "Benchmark graphs for testing community detection algorithms," *Phys. Rev. E*, vol. 78, p. 046110, Oct 2008.
 [23] S. Fortunato and M. Barthélemy, "Resolution limit in community detection," *Proceedings of the National Academy of Sciences*, vol. 104, no. 1, pp. 36–41, 2007.