

# Parallel Algorithms for Switching Edges in Heterogeneous Graphs<sup>☆</sup>

Hasanuzzaman Bhuiyan<sup>a,b,\*</sup>, Maleq Khan<sup>b,\*\*</sup>, Jiangzhuo Chen<sup>b</sup>, Madhav  
Marathe<sup>a,b</sup>

<sup>a</sup>*Department of Computer Science, Virginia Tech,  
2202 Kraft Drive, Blacksburg, VA 24061, USA*

<sup>b</sup>*Network Dynamics and Simulation Science Laboratory, Virginia Bioinformatics Institute,  
Virginia Tech, 1880 Pratt Drive, Blacksburg, VA 24061, USA*

---

---

---

<sup>☆</sup>A preliminary version of this paper entitled “Fast Parallel Algorithms for Edge-Switching to Achieve a Target Visit Rate in Heterogeneous Graphs” [1] has appeared in the proceedings of the 43rd International Conference on Parallel Processing (ICPP), 2014. This work has been partially supported by DTRA CNIMS Contract HDTRA1-11-D-0016-0001, DTRA Grant HDTRA1-11-1-0016, NSF NetSE Grant CNS-1011769, NSF SDCI Grant OCI-1032677, NIH MIDAS Grant 5U01GM070694-11 and NSF DIBBs Grant ACI-1443054.

\*Corresponding author

\*\*Principal corresponding author

*Email addresses:* [mhb@vbi.vt.edu](mailto:mhb@vbi.vt.edu) (Hasanuzzaman Bhuiyan), [maleq@vbi.vt.edu](mailto:maleq@vbi.vt.edu) (Maleq Khan), [chenj@vbi.vt.edu](mailto:chenj@vbi.vt.edu) (Jiangzhuo Chen), [mmarathe@vbi.vt.edu](mailto:mmarathe@vbi.vt.edu) (Madhav Marathe)

## Abstract

An edge switch is an operation on a graph (or network) where two edges are selected randomly and one of their end vertices are swapped with each other. Edge switch operations have important applications in graph theory and network analysis, such as in generating random networks with a given degree sequence, modeling and analyzing dynamic networks, and in studying various dynamic phenomena over a network. The recent growth of real-world networks motivates the need for efficient parallel algorithms. The dependencies among successive edge switch operations and the requirement of keeping the graph simple (i.e., no self-loops or parallel edges) as the edges are switched lead to significant challenges in designing a parallel algorithm. Addressing these challenges requires complex synchronization and communication among the processors leading to difficulties in achieving a good speedup by parallelization. In this paper, we present distributed memory parallel algorithms for switching edges in massive networks. These algorithms provide good speedup and scale well to a large number of processors. On a network with 20 million vertices and 587 million edges, a speedup of 110 is achieved using 640 processors. One of the steps in our edge switch algorithms requires the computation of multinomial random variables in parallel. This paper presents the first non-trivial parallel algorithm for the problem, which achieves a speedup of 925 using 1024 processors.

*Keywords:* edge switch, random network generation, network dynamics, multinomial distribution, parallel algorithms.

## 1. Introduction

Edge switch, also known as edge swap, edge flip, edge shuffle, edge rewiring, etc., is an operation that swaps the end vertices of the edges in a network. Many variations of this problem have been studied [2, 3, 4, 5, 6, 7, 8, 9, 10] with diverse real-world applications. In the most commonly used edge switch operation, two randomly selected edges  $(a, b)$  and  $(c, d)$  are replaced with edges  $(a, d)$  and  $(c, b)$  respectively, i.e., the end vertices of the selected edges are swapped with each other. This operation is repeated either a given number of times or until a specified criteria is satisfied. It is easy to see that an edge switch operation preserves the degree of each vertex.

This problem has many important applications. It can be used in generating random networks with a given degree sequence. There has been a lot of work on random graph generation, because of the popularity of network models in diverse applications. Most of the prior work involves sequential algorithms, and much of it is restricted to regular graphs; we briefly summarize the main approaches here. A popular method for random graph generation is the *configuration model* (also referred to as the “pairing” model) [11, 12, 13], which involves creating stubs for vertices, choosing pairs of stubs at random, and then connecting them by edges. Unfortunately, this leads to parallel edges, unless the degrees are very small. This basic approach has been modified in various ways to avoid parallel edges in the case of regular graphs [13, 14, 15] (see [12] for a good discussion). Blitzstein et al. [12] gives a simple algorithm for generating random graphs with a given degree sequence using sequential importance sampling, based on the Erdős-Gallai characterization.

By using the Havel-Hakimi method [16], a network can be generated following a given degree sequence. Since it is deterministic, this method generates the same network each time it is run with the same degree sequence whereas there can be many different networks with the same degree distribution. However, edge switching can be combined with the Havel-Hakimi method to generate a random network with a given degree sequence [4, 2, 3]. Once a network is gen-

erated using the Havel-Hakimi method, by randomly switching the edges we can generate a random network with the same degree sequence. The mixing time was shown to be bounded by a large polynomial by Cooper et al. [2], and extended by Feder et al. [3] to variants of the edge switch process.

Edge switching is also used in modeling and studying various dynamic networks such as peer-to-peer networks [3]. Other applications of edge switching include the generation of randomly labeled bipartite graphs with a given degree sequence [6], independent realizations of graphs with a prescribed joint degree distribution using a Markov chain Monte Carlo approach [7], and studying the sensitivity of network topology on dynamics over a network such as disease dynamics over a social contact network [17].

Edge switching can be paired with additional constraints such as imposing a connectivity requirement, allowing or not allowing parallel edges and loops, etc. NetworkX [18] has a sequential implementation of edge switching that does not allow parallel edges, but allows loops, and provides the option of imposing connectivity constraints on the graph. A connectivity constraint requires a graph to remain connected after an edge switch operation. Some theoretical studies of edge switching for restricted graph classes can be found in the literature, such as the study of mixing time of the Markov chain introduced by this operation [2, 4]. However, no effort was given to design parallel algorithms for switching edges in a graph. For smaller graphs, sequential implementation of edge switching suffices, but this may not work for massive networks for the following reasons: (i) a massive network with billions of edges simply may not fit in the memory of a single computing machine, and (ii) a sequential algorithm may take a prohibitively long time. These issues can be addressed by a distributed memory parallel algorithm where the network is partitioned and each processor contains one partition.

**Our Contributions.** In this paper, we present distributed memory parallel algorithms for switching edges in massive graphs with the constraint that the graph remains simple. The dependencies among successive edge switch operations and the requirement of keeping the graph simple lead to significant chal-

lenges in designing a parallel algorithm. To deal with these challenges, it requires complex synchronization and communication among the processors, which in turn makes it challenging to gain any speedup by parallelization. The performance of the algorithms also depend on partitioning of the graph. We study several partitioning schemes in conjunction with the algorithms and present their trade-offs. One of the parallel algorithms achieves a speedup of 110 using 640 processors, allowing it to perform 115 billion edge switches in a very large power-law network with 10 billion edges in less than 3 hours using 1024 processors. The algorithms require generating multinomial random variables in parallel, which is also a non-trivial problem. To the best of our knowledge, there is no existing parallel algorithm for the problem, and we present here a novel parallel algorithm for generating multinomial random variables, which achieves a speedup of 925 using 1024 processors.

**Organization.** The rest of the paper is organized as follows. Section 2 describes the preliminaries and notations used in the paper. The edge switching problem and the sequential algorithm are briefly explained in Section 3. We present our main parallel algorithm for switching edges in Section 4. In Section 5, we present our parallel algorithms using several hash-based partitioning schemes. The parallel algorithm for generating multinomial random variables is presented in Section 6. Finally, we conclude in Section 7.

## 2. Preliminaries

Below are the notations, definitions and computation model used in this paper.

**Notations.** We are given a simple graph  $G = (V, E)$ , where  $V$  is the set of vertices, and  $E$  is the set of edges. A *simple graph* is an undirected graph with no self-loops or parallel edges. A *self-loop* is an edge from a vertex to itself. *Parallel edges* are two or more edges connecting the same pair of vertices. There are total  $n = |V|$  vertices labeled as  $0, 1, 2, \dots, n - 1$ , and  $m = |E|$  edges in the graph  $G$ . If  $(u, v) \in E$ , we say  $u$  and  $v$  are *neighbors* of each other. The

neighbors of a vertex  $u \in V$  are stored in the *adjacency list* of  $u$ , denoted as  $N(u)$ , i.e.,  $N(u) = \{v \in V | (u, v) \in E\}$ . The degree of  $u$  is  $d_u = |N(u)|$ . The terms *node* and *vertex*, *graph* and *network*, *neighbor list* and *adjacency list*, *loop* and *self-loop*, *label* and *vertex-id* are used interchangeably throughout the paper. We use  $H, K, M$  and  $B$  to denote hundreds, thousands, millions and billions, respectively; e.g.,  $1M$  stands for one million. For the parallel algorithms, let  $p$  be the number of processors, and  $P_i$  the processor with rank  $i$ .

**Edge Switch.** An edge switch operation replaces two edges  $e_1 = (u_1, v_1)$  and  $e_2 = (u_2, v_2)$ , selected uniformly at random from  $E$ , by new edges  $e_3 = (u_1, v_2)$  and  $e_4 = (u_2, v_1)$ , as shown in Figure 1. If  $u_1 = v_2$  or  $u_2 = v_1$ , then the above edge switch creates self-loops. The edge switch creates parallel edges, if edge  $(u_1, v_2)$  or  $(u_2, v_1)$  already exist in the graph.

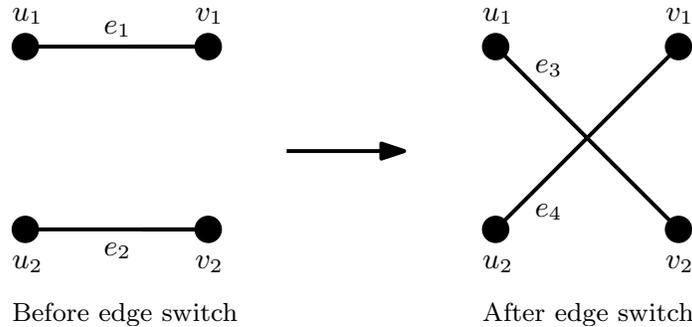


Figure 1: An edge switch operation replaces two randomly selected edges  $e_1 = (u_1, v_1)$  and  $e_2 = (u_2, v_2)$  by new edges  $e_3 = (u_1, v_2)$  and  $e_4 = (u_2, v_1)$ .

**Visit Rate.** Due to edge switch operations, some edges of the given graph  $G$  are changed (visited), and some edges that do not participate in any edge switch operation remain unchanged (not visited). We define *visit rate* as the fraction of edges of  $G$  that have been changed by a sequence of edge switch operations. If  $m'$  is the number of edges of  $G$  that have been changed due to switching edges, then visit rate  $x = m'/m$ .

**Binomial Distribution.** Suppose that  $N$  independent trials are to be performed, where each trial results in a success with probability  $q$ , and in a failure with probability  $(1 - q)$ . If  $X$  represents the number of successes that

occur among  $N$  trials, then  $X$  is said to be a binomial random variable. The distribution of  $X$  is a binomial distribution with parameters  $N$  and  $q$ , and denoted by equation (1). The probability of getting exactly  $i$  successes in  $N$  trials is given in equation (2).

$$X \sim \mathcal{B}(N, q) \tag{1}$$

$$\Pr\{X = i\} = \binom{N}{i} q^i (1 - q)^{N-i} \tag{2}$$

**Multinomial Distribution.** Let  $N$  be the number of independent trials to be performed, where each trial has  $\ell$  possible outcomes  $0, 1, \dots, \ell - 1$  with probability  $q_0, q_1, \dots, q_{\ell-1}$  respectively, such that  $q_i \geq 0$  for  $0 \leq i \leq \ell - 1$  and  $\sum_i q_i = 1$ . Let the random variable  $X_i$  indicates the number of times the outcome  $i$  appears among  $N$  independent trials. Then  $X = \langle X_0, X_1, \dots, X_{\ell-1} \rangle$  has a multinomial distribution with parameters  $N, q_0, q_1, \dots, q_{\ell-1}$ , and denoted as follows.

$$\langle X_0, X_1, \dots, X_{\ell-1} \rangle \sim \mathcal{M}(N, q_0, q_1, \dots, q_{\ell-1}) \tag{3}$$

**Computation Model.** We develop algorithms for distributed memory parallel systems. Each processor has its own local memory. The processors do not have any shared memory and can communicate with each other and exchange data by message passing.

### 3. Sequential Edge Switch

We are given a simple graph  $G = (V, E)$  and the number of edge switch operations  $t$  to be performed. A random edge switch operation comprises of choosing a pair of edges  $e$  and  $e'$  uniformly at random from the set of edges in the graph and switching the end vertices of  $e$  and  $e'$ . A sequence of such  $t$  pairs of edges are switched.

#### 3.1. Determining the Number of Edges to Switch for a Given Visit Rate

During an edge switch operation, a selected edge can be categorized as one of the following two types. (i) *Original edge*: an edge that has not participated in

any of the previous edge switch operations and is still unchanged. *(ii) Modified edge:* any edge participating in an edge switch operation is replaced by a new edge, and such a new edge is called a modified edge.

Let  $T$  be the total number of edges switched to achieve a visit rate  $x$ . Since edge switching is a random process, performing the same number of edge switch operations in different executions of the same edge switching algorithm may exhibit different visit rates. Thus having an exact value of  $T$  in advance is not possible. However, we can calculate the *expected* value of  $T$  as described below. As we demonstrate later in this section, using this expected value of  $T$  leads to a very close approximation of the visit rate. Finding the expected value of  $T$  is similar to the coupon collector problem [19]. Our goal is to have  $m' = mx$  modified edges in the graph by switching a sequence of pairs of edges. The remainder  $(m - m')$  of the edges remain unchanged. At some point there are already  $(i - 1)$  modified edges in the graph. From this point to have the  $i$ -th modified edge we need  $T_i$  number of edges switched. The probability of selecting the  $i$ -th original edge from the graph, given that there are  $(i - 1)$  modified edges, is  $p_i = \frac{m-(i-1)}{m}$ . Here,  $T$  and  $T_i$  are random variables, and  $T_i$  has geometric distribution with expectation  $1/p_i$ . Using the linearity of expectation,

$$\begin{aligned}
 E[T] &= \sum_{i=1}^{mx} E[T_i] = \sum_{i=1}^{mx} \frac{1}{p_i} = \sum_{i=1}^{mx} \frac{m}{m - (i - 1)} \\
 &= m \left( \sum_{i=1}^m \frac{1}{i} - \sum_{i=1}^{m(1-x)} \frac{1}{i} \right) \\
 &= m (H_m - H_{m(1-x)})
 \end{aligned} \tag{4}$$

where  $H_m$  is the  $m$ -th harmonic number. For large  $m$ ,  $H_m \approx \ln m$ , and consequently  $E[T] \approx -m \ln(1 - x)$  for  $x < 1$ , and  $E[T] \approx m \ln m$  for  $x = 1$ . Note that every edge switch operation involves two edges. Now if we assign  $t$  to be  $E[T]/2$ , we obtain a visit rate extremely close to  $x$  as demonstrated below.

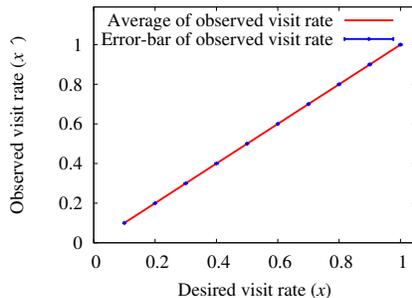


Figure 2: Observed visit rate is almost equal to the desired visit rate for Miami network. The error is so small that the error-bar is almost invisible.

Desired visit rate	Observed visit rate	
	Average error rate (%)	Standard deviation
0.1	0.00745	8.13E-6
0.2	0.00858	1.41E-5
0.3	0.00907	1.76E-5
0.4	0.00802	3.52E-5
0.5	0.00687	2.34E-5
0.6	0.00650	3.38E-5
0.7	0.00701	4.37E-5
0.8	0.01030	5.55E-5
0.9	0.00824	4.46E-5
1.0	2.4E-6	2.06E-8

Table 1: Average error rate and standard deviation of observed visit rates for Miami network are near to 0. For each desired visit rate, 10 experiments are performed.

We perform experiments on a contact network of Miami city having  $m = 52.7M$  edges (see Section 4.7 for details) to achieve a visit rate of  $x = 1$ , i.e., visit all of the  $52.7M$  edges. The expected value of  $T$  is calculated using  $E[T] \approx m \ln m$ , and the edge switching algorithm performs  $t = 468.5M$  edge switch operations. We repeat this experiment 10 times and observe a visit rate of  $x' = 1$  (visiting all edges) for 20% times,  $x' = 0.99999998$  (visiting all but one edge) for 60% times and  $x' = 0.99999994$  (visiting all but three edges) for 20% times. Thus the observed visit rates are extremely close to  $x$ . We perform additional experiments for desired visit rates  $x = 0.1, 0.2, \dots, 1$  on Miami network. Each experiment is repeated 10 times. Figure 2 demonstrates that the observed visit rates are almost equal to the desired visit rates. We plot the minimum and maximum of observed visit rates using error-bars. These values are so close to the desired visit rates that they almost overlap with each other and it is difficult to distinguish them in the figure. To better understand the differences between the desired and observed visit rates, we further compute the average error rate and standard deviation of the observed visit rates, which are shown in Table 1. The average error rate (%) is calculated as  $\frac{\sum_i |x_i - x'_i|}{ex_i} \times 100\%$ , where  $x_i$  and  $x'_i$  are

the desired and observed visit rates, respectively, in the  $i$ -th experiment and  $e$  is the total number of experiments. The maximum, minimum and average error rates of the total 100 experiments are 0.027%, 0% and 0.007%, respectively, which are almost negligible. Therefore, for large  $m$ , we achieve a very close approximation of  $x$ , which is sufficient for almost all practical purposes.

Note that we can mark the modified edges and always select two original edges for the next edge switch operation. In such a case for a visit rate  $x$  to have  $mx$  modified edges, we simply need to perform  $mx/2$  edge switch operations. For a specific application, one can do so. If we do not allow a modified edge to participate in any later edge switch operation, the process may not produce many networks with the same degree sequence. Unrestricted and independent random choice of the edges helps us obtain a random graph from the space of the graphs with the same degree sequence.

Furthermore, visit rate can also be defined in other ways and converted to  $t$ . Our parallel algorithms can be used to perform  $t$  edge switch operations, irrespective of how  $t$  is obtained.

### 3.2. Keeping the Graph Simple

Because the edge switching problem deals with a simple graph, we need to ensure that none of the edge switch operations create self-loops or parallel edges. An edge switch between edges  $(u_1, v_1)$  and  $(u_2, v_2)$  creates

- **Parallel edge:** if  $u_1 \in N(v_2)$ ,  $v_2 \in N(u_1)$ ,  $u_2 \in N(v_1)$  or  $v_1 \in N(u_2)$ .
- **Self-loop:** if  $u_1 = v_2$  or  $u_2 = v_1$ .

An edge switch operation does not make any change to the graph if the pair of edges remain the same after switching, and we say such an edge switch operation is *useless*. An edge switch between  $(u_1, v_1)$  and  $(u_2, v_2)$  is useless if  $u_1 = u_2$  or  $v_1 = v_2$ . For an edge switch operation, two edges are selected and switched if the switch is not useless and does not create parallel edges or loops.

### 3.3. Switching Edges Sequentially

The sequential algorithm is quite simple. Select a pair of edges uniformly at random and switch them if the resultant graph remains simple. This operation is repeated until  $t$  pairs of edges are switched. The graph, specifically the edge set, dynamically changes with the course of edge switch process. Let  $G' = (V, E')$  be such a graph where  $E'$  is the current set of edges at a given time. Algorithm 1 shows the pseudocode of switching edges sequentially. The adjacency list of a vertex can be stored using a balanced binary tree. Searching such an adjacency list of a vertex  $u$  to determine the possibility of parallel edge creation takes  $O(\log d_u)$  time. If  $(u_1, v_1)$  and  $(u_2, v_2)$  are the edges participating in the  $i$ -th edge switch operation, then the time to switch  $t$  pairs of edges is  $O\left(\sum_{i=1}^t \sum_{j \in \{u_1, v_1, u_2, v_2\}} \log d_j\right) \leq O(t \log d_{max})$ , where  $d_{max}$  is the maximum degree of a vertex in the graph. Note that if an edge switch operation attempts to create a parallel edge or a loop, or is useless, the edge switch operation is restarted by selecting a new pair of edges. For a large and relatively sparse network, this probability is very small. As a result, the number of edge switch operations restarted is significantly smaller than  $t$ . Thus we have the runtime  $O(t \log d_{max})$ .

---

#### Algorithm 1 SEQUENTIAL EDGE SWITCH $(G, t)$

---

- 1: **for**  $i = 1$  to  $t$  **do**
  - 2:    $(u_1, v_1), (u_2, v_2) \leftarrow$  two uniform random edges in  $E'$
  - 3:   **if**  $u_1 = u_2, v_1 = v_2, u_1 = v_2, u_2 = v_1, u_1 \in N(v_2),$  or  $u_2 \in N(v_1)$  **then**
  - 4:     go to line 4
  - 5:   Replace  $(u_1, v_1)$  and  $(u_2, v_2)$  by  $(u_1, v_2)$  and  $(u_2, v_1)$  respectively
- 

## 4. Parallel Edge Switch

Although the sequential algorithm is very simple, parallelizing the simple edge switch operations turns out to be a non-trivial problem for the following reasons.

- Multiple pairs of edges are selected and switched simultaneously by different processors in the parallel process, whereas the sequential process selects and switches a sequence of pairs of edges, one pair after another. Designing a parallel algorithm by maintaining an equivalent stochastic process as of the sequential one leads to significant challenges.
- The requirement of keeping the graph simple requires complex synchronization and communication among the processors. To achieve a good speedup by parallelization, we need to design an efficient algorithm by minimizing such communication and computation costs.

In this section, we present an efficient parallel algorithm for switching edges in massive graphs, accompanied by a rigorous comparative study of several parallel algorithms using various partitioning schemes later in Section 5.

#### 4.1. Overview of the Algorithm

The input graph  $G$  is partitioned and distributed among the  $p$  processors. Each partition contains a subset of the edges, and is assigned to a processor. All the processors then perform  $t$  edge switch operations in parallel. We need to consider two cases for an edge switch operation.

- **Local Switch.** Both edges may be selected from the same partition (or processor), and this is referred to as *local switch*.
- **Global Switch.** The edges may be selected from different partitions, which is referred to as *global switch*.

The processors may need to communicate with each other in order to complete the edge switch operation. We explain the details below in the following order: (i) data structures, (ii) partitioning the network, (iii) switching a pair of edges by a single processor, (iv) simultaneous edge switches by all processors, (v) properties of parallel edge switch, and (vi) experimental results.

#### 4.2. Data Structures

A graph can be stored as adjacency lists or an adjacency matrix. In an adjacency matrix, the existence of any edge can be determined in constant time, however it takes  $O(n^2)$  memory. Our algorithms use adjacency lists which require  $O(m + n)$  memory. Usually,  $N(u)$  contains all neighbors of  $u$ .

**Reduced Adjacency List.** For an edge  $(u, v)$ , if  $N(u)$  and  $N(v)$  belong to different partitions, the edge can be selected from two different partitions and participate in two different edge switch operations at the same time leading to an inconsistency. To ensure that any edge  $(u, v)$  can be selected only from one partition, only the neighbors with higher labels are kept in the adjacency list of a vertex  $u$ , i.e.,  $N(u) = \{v \in V | (u, v) \in E, u < v\}$ , which is referred to as *reduced adjacency list*. Although it is possible to deal with the above issue by keeping all neighbors in the adjacency list, it will incur more communication costs. Every edge switch operation involves updating four vertices' adjacency lists: one update for each end vertex of an edge. A reduced adjacency list minimizes the number of updates to only two or three vertices' adjacency lists; the details are discussed later in Section 4.4. Thus a reduced adjacency list reduces memory footprint, and communication and computation costs.

**Straight and Cross Edge Switch.** A difficulty arises from using the reduced adjacency list. If  $N(u)$  contains all the neighbors of  $u$ , any edge  $(u_1, v_1)$  can be selected either as  $(u_1, v_1)$  from  $N(u_1)$  (considering ordered pair), or as  $(v_1, u_1)$  from  $N(v_1)$ . The probability of being selected each way is  $1/2m$ . Let  $(u_1, v_1)$  and  $(u_2, v_2)$  (considering no ordering) be two edges selected for an edge switch operation. Depending on whether the edge  $(u_1, v_1)$  is selected from  $N(u_1)$  or  $N(v_1)$ , and the other edge  $(u_2, v_2)$  is chosen from  $N(u_2)$  or  $N(v_2)$ , the edges are replaced by either  $(u_1, v_2)$  and  $(u_2, v_1)$ , or  $(u_1, u_2)$  and  $(v_1, v_2)$ . Assuming  $u_1 < v_1$  and  $u_2 < v_2$ , there is no possibility of selecting edges as  $(v_1, u_1)$  and  $(v_2, u_2)$  (considering ordered pair) due to the use of a reduced adjacency list. Therefore, an edge switch between  $(u_1, v_1)$  and  $(u_2, v_2)$  (considering unordered pair) misses the chance of generating the edges  $(u_1, u_2)$  and  $(v_1, v_2)$ . This problem is solved by replacing the selected edges by either  $(u_1, u_2)$  and  $(v_1, v_2)$  with

probability  $1/2$ , referred to as *straight switch*, or  $(u_1, v_2)$  and  $(u_2, v_1)$  with probability  $1/2$ , referred to as *cross switch*, as shown in Figure 3.

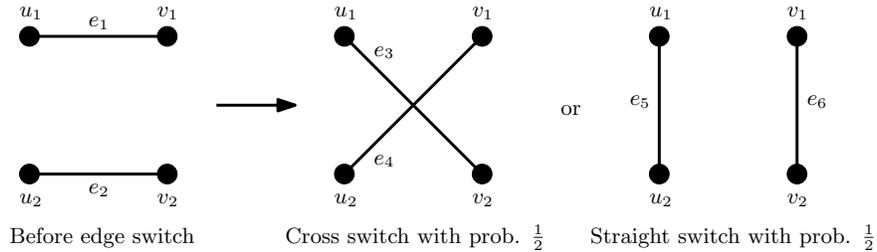


Figure 3: Straight and cross edge switch.

#### 4.3. Partitioning the Network

For a given simple graph  $G = (V, E)$ , we partition  $V$  into  $p$  disjoint subsets,  $V_0, V_1, \dots, V_{p-1}$ , such that  $\bigcup_i V_i = V$ . The reduced adjacency list of a vertex entirely belongs to one partition, and a subset of consecutive (in terms of vertex labels) vertices are assigned to each partition. The input graph is partitioned in such a way that each partition contains roughly  $m/p$  edges. Let  $V_i$  be the subset of vertices (having consecutive vertex labels) and  $E_i$  be the subset of edges in the partition belonging to  $P_i$  (processor with rank  $i$ ) such that  $E_i = \{(u, v) \in E \mid u \in V_i, u < v\}$ . Note that the partitions are disjoint, i.e.,  $E_i \cap E_j = \emptyset$  for  $i \neq j$ , and  $\bigcup_i E_i = E$ . As the graph is dynamically changing with edge switch operations, let us denote  $E_i$  to be the current set of edges in  $P_i$  at a given time. We refer to this partitioning scheme as ***Consecutive Partitioning (CP)***.

#### 4.4. Switching a Pair of Edges by a Single Processor

A simple approach to perform an edge switch operation is that processor  $P_i$  can select one pair of edges uniformly at random from the entire graph (i.e., selecting two processors from  $[0, p - 1]$  and request them for edges) and switch them by exchanging messages among the processors. However, this approach incurs significant synchronization and communication costs. Instead,  $P_i$  selects one edge  $(u_1, v_1)$ , referred to as *first edge*, uniformly at random from  $E_i$ , and

another edge  $(u_2, v_2)$ , referred to as *second edge*, from the entire graph. To select a second edge,  $P_i$  selects a processor  $P_j$  with probability  $|E_j|/|E|$  and requests  $P_j$  to select an edge  $(u_2, v_2)$  from  $E_j$  uniformly at random. If  $P_i = P_j$ , then it is a *local switch*, otherwise it is a *global switch*. Due to the use of reduced adjacency lists, one of the replacing edges ( $e_3, e_4, e_5$  or  $e_6$  in Figure 3) may belong to a different processor  $P_k$  ( $P_i \neq P_k \neq P_j$ ); in this case, processor  $P_i, P_j$  and  $P_k$  work together to update the reduced adjacency lists of respective vertices by exchanging messages and thus complete the edge switch operation. A high-level overview of an edge switch operation is given in Algorithm 2. During the course of an edge switch operation, if any processor  $P_x$  detects a possibility of creating loops or parallel edges,  $P_x$  notifies all other processors that are involved in the edge switch operation. Then the initiating processor ( $P_i$  in the above example) restarts the edge switch operation by selecting a new pair of edges.

---

**Algorithm 2** SWITCHING A PAIR OF EDGES INITIATED BY  $P_i$

---

- 1:  $e_1 \leftarrow$  a uniform random edge in  $E_i$
  - 2:  $P_j \leftarrow$  a random processor in  $[0, p - 1]$ , where probability of choosing  $P_x$  is  $\frac{|E_x|}{|E|}$
  - 3: **if**  $P_i = P_j$  **then**
  - 4:   Choose an edge  $e_2$  from  $E_i$  to switch with edge  $e_1$
  - 5:   Switch the edges  $e_1$  and  $e_2$  ( $P_i$  may communicate with a different processor  $P_k$  to complete the edge switch operation)
  - 6: **else**
  - 7:   Send message  $\langle e_1, \mathbf{request\ to\ select\ an\ edge\ from\ } E_j \rangle$  to  $P_j$
  - 8:   Upon receipt of the above message,  $P_j$  executes the following:
  - 9:   Choose an edge  $e_2$  from  $E_j$  to switch with edge  $e_1$
  - 10:  $P_i$  and  $P_j$  work together to switch  $e_1$  and  $e_2$  ( $P_j$  may communicate with a different processor  $P_k$  to complete the edge switch operation)
- 

**Local Switch.**  $P_i$  selects two edges  $(u_1, v_1)$  and  $(u_2, v_2)$  from  $E_i$  uniformly

at random such that the edge switch does not create loops, and it is not useless.  $P_i$  decides between a *straight* and a *cross switch* with equal probability. If it is a *cross switch*,  $P_i$  checks whether  $(u_1, v_2)$  and  $(u_2, v_1)$  create parallel edges. If no parallel edge is created,  $P_i$  removes  $(u_1, v_1)$  and  $(u_2, v_2)$ , adds  $(u_1, v_2)$  and  $(u_2, v_1)$ , thus completing the edge switch operation. If the edge switch is a *straight switch*,  $P_i$  determines  $P_k$  such that  $\min(v_1, v_2) \in V_k$ . If  $P_i = P_k$ ,  $P_i$  determines whether  $(u_1, u_2)$  and  $(v_1, v_2)$  create parallel edges. If they do not create any parallel edge,  $P_i$  removes  $(u_1, v_1)$  and  $(u_2, v_2)$ , adds  $(u_1, u_2)$  and  $(v_1, v_2)$  and completes the edge switch operation. If  $P_i \neq P_k$ ,  $P_i$  checks whether  $(u_1, u_2)$  creates parallel edges. If the graph remains simple,  $P_i$  sends a message to  $P_k$  requesting to add  $(v_1, v_2)$ . If  $(v_1, v_2)$  does not create parallel edges,  $P_k$  adds  $(v_1, v_2)$  and sends a message back to  $P_i$  informing it of the updates at  $P_k$ . Upon receiving this message,  $P_i$  removes  $(u_1, v_1)$ ,  $(u_2, v_2)$  and adds  $(u_1, u_2)$ .

**Global Switch.** In a *global switch*, two edges are selected from two different processors, say  $P_i$  and  $P_j$ ,  $i < j$ . Assuming  $P_i$  initiates the edge switch operation,  $P_i$  selects an edge  $e_1 = (u_1, v_1)$  from  $E_i$  uniformly at random.  $P_i$  sends a message containing the edge  $e_1$  and a request to select an edge from  $E_j$ , to  $P_j$ . Upon receiving this message from  $P_i$ , processor  $P_j$  selects  $e_2 = (u_2, v_2)$  from  $E_j$  uniformly at random, and decides between a *straight* and a *cross switch* with equal probability. At this point,  $P_j$  knows the new edges that will replace  $e_1$  and  $e_2$ ; we refer to these new edges as *potential edges* until the updates take place. Next we describe the *cross switch* in detail.

Processor  $P_j$  checks whether  $u_2 = v_1$  and  $v_1 = v_2$  to detect a loop and a useless edge switch respectively. If it does not create a loop and is not useless,  $P_j$  determines  $P_k$  such that  $\min(u_2, v_1) \in V_k$ . We need to consider the following three cases.

i. *Case  $P_k = P_j$ :*

$P_j$  checks whether  $(u_2, v_1)$  creates parallel edges. If a parallel edge is not created, then  $P_j$  sends  $v_2$  to  $P_i$ .  $P_i$  checks whether  $(u_1, v_2)$  creates parallel edges. If the graph remains simple,  $P_i$  removes edge  $(u_1, v_1)$ , adds

edge  $(u_1, v_2)$ , and sends a message back to  $P_j$  informing the updates at  $P_i$ . Upon receiving this message,  $P_j$  removes  $(u_2, v_2)$  and adds  $(u_2, v_1)$ , thus completing the edge switch operation.

ii. *Case  $P_k = P_i$ :*

$P_j$  sends a message, containing  $e_2$  and a request to add both the new edges to  $P_i$ . Processor  $P_i$  checks whether  $(u_1, v_2)$  and  $(u_2, v_1)$  create parallel edges. If no parallel edge is created,  $P_i$  removes  $(u_1, v_1)$ , adds edges  $(u_1, v_2)$  and  $(u_2, v_1)$ , and sends a message back to  $P_j$  indicating the updates at  $P_i$ . Then  $P_j$  completes the edge switch operation by removing  $(u_2, v_2)$ .

iii. *Case  $P_i \neq P_k \neq P_j$ :*

$P_j$  sends  $(u_2, v_1)$  and  $v_2$  to  $P_k$ . If  $(u_2, v_1)$  does not create any parallel edge,  $P_k$  sends  $v_2$  to  $P_i$ .  $P_i$  checks whether  $(u_1, v_2)$  creates any parallel edge. If the graph remains simple,  $P_i$  removes  $(u_1, v_1)$ , adds  $(u_1, v_2)$ , and sends messages to  $P_j$  and  $P_k$  notifying the updates taken place at  $P_i$ . Then  $P_j$  removes edge  $(u_2, v_2)$ , and  $P_k$  adds edge  $(u_2, v_1)$ , thus completing the edge switch operation.

A similar approach is followed for  $i > j$  and for a straight switch as well. The use of reduced adjacency lists eliminates the following two constraints: (i)  $u_1 = u_2$ , and (ii)  $u_1 = v_2$  if  $i < j$ , or  $u_2 = v_1$  if  $i > j$ .

#### 4.5. Simultaneous Edge Switches by All Processors

In a sequential algorithm, pairs of edges are selected randomly, one pair after another; as a result, the number of edges selected from each partition  $E_i$  may not be equal. To have an equivalent parallel algorithm, we need to select the same number of edges from each partition  $E_i$  as the sequential algorithm would do. Let  $X_i$  be the number of first edges selected from  $E_i$  by a sequential algorithm. A sequential algorithm does not need to know  $X_i$  in advance. However, for the parallel algorithm, for each  $i$ ,  $X_i$  needs to be determined in advance so that processors can simultaneously perform edge switches in parallel. For any edge switch operation, the probability that the first edge is selected from  $E_i$  is

$q_i = |E_i|/|E|$  for  $i = 0, 1, \dots, p-1$ , and we have  $\sum_{i=0}^{p-1} q_i = 1$ . Then it is easy to see that the random variables  $X_i$  for  $i = 0, 1, \dots, p-1$  are multinomially distributed with parameters  $(t, q_0, q_1, \dots, q_{p-1})$ ; i.e.,

$$\langle X_0, X_1, \dots, X_{p-1} \rangle \sim \mathcal{M}(t, q_0, q_1, \dots, q_{p-1}) \quad (5)$$

The time complexity of the best known sequential algorithm, known as *conditional distributed method* [20], for generating multinomial variables is  $\Theta(t)$ . Thus to have an efficient parallel algorithm for our edge switching problem, we need to use an efficient parallel algorithm for generating multinomial random variables. To the best of our knowledge, there is no existing parallel algorithm for this problem. In Section 6, we present an efficient parallel algorithm for computing multinomial random variables that runs in  $O\left(\frac{t}{p} + p \log p\right)$  time.

Each processor  $P_i$  simultaneously performs  $X_i$  number of edge switches and serves other processors' requests as well. After completing one edge switch,  $P_i$  proceeds to its next edge switch operation. Below we discuss two *issues* that arise from performing edge switch operations simultaneously.

**1. Creating parallel edges in a new way.** Even after maintaining all the constraints to keep a graph simple, parallel edges can be created in a different way. As multiple pairs of edges are switched by multiple processors simultaneously, the same new edge can be created by multiple processors at the same time. For example, more than one instance of an edge  $(u, v)$  is created simultaneously if more than one of the following four edge switches are performed simultaneously by different processors, where ‘ $-$ ’ denotes an end vertex of an edge. (i) Cross edge switch between  $(u, -)$  and  $(-, v)$ . (ii) Cross edge switch between  $(-, u)$  and  $(v, -)$ . (iii) Straight edge switch between  $(u, -)$  and  $(v, -)$ . (iv) Straight edge switch between  $(-, u)$  and  $(-, v)$ . Keeping track of *potential edges* at each processor ensures no parallel edges will be created in the above mentioned way.

**2. Changing probability values with the course of edge switch process.** As the edges are switched, the number of edges changes (i.e., increases or decreases) among the partitions due to the use of reduced adjacency lists.

As a result, the probability values ( $q_i$ ) of selecting edges from different partitions change, which need to be updated dynamically. However, updating the probability values after every edge switch operation incurs a large communication costs, which in turn slows down the algorithm significantly. To deal with this difficulty, the processors perform a fixed number of edge switch operations, referred to as *step-size* and denoted by  $s$ , in a step, and then update the probability values that are used in the next step. Therefore, the algorithm performs edge switch operations in a number of steps. At the beginning of each step,  $s$  edge switch operations are distributed among  $p$  processors using multinomial distribution. The program terminates when all of the  $t$  edge switch operations are performed in  $\lceil \frac{t}{s} \rceil$  steps. With a reasonable step-size, a very close approximation of the sequential algorithm is achieved. The experimental results are shown later in Section 4.7.

**Summary of the Parallel Algorithm.** Let  $s$  be the *step-size*, and  $q$  be the probability vector  $\langle q_0, q_1, \dots, q_{p-1} \rangle$ . All the processors perform  $s$  edge switch operations in one step, thus requiring total  $\lceil \frac{t}{s} \rceil$  number of steps. If  $t \% s \neq 0$ ,  $(t - s \lfloor \frac{t}{s} \rfloor)$  number of edge switch operations are performed in the last step. Below is a summary of the parallel algorithm.

1) **Generating multinomial random variables.** At the beginning of each step,  $s$  edge switch operations are distributed among  $p$  processors using the parallel algorithm for generating multinomial random variables with parameters  $(s, q_0, q_1, \dots, q_{p-1})$ . This takes  $O\left(\frac{s}{p} + p \log p\right)$  time. Let us denote  $S_i$  to be the number of edge switch operations that a processor  $P_i$  performs in the current step.

2) **Performing edge switch operations.** To perform an edge switch operation, a processor  $P_i$  selects one edge  $e_1$  from  $E_i$ , and the other edge  $e_2$  from the entire graph, and completes the edge switch operation in conjunction with other processors (details in Section 4.4). Each processor  $P_i$  simultaneously performs such  $S_i$  number of edge switch operations and serves other processors' requests as well. For an edge switch operation, a constant amount of message exchange is required; edges are updated in constant time and checking for parallel edges

takes  $O(\log d_{max})$  time. Thus, performing  $S_i$  edge switch operations at  $P_i$  takes  $O(S_i \log d_{max})$  time.

**3) Updating probability vector and termination.** After completing  $S_i$  edge switch operations in the current step,  $P_i$  sends *end-of-step* signals (or messages) to each processor requiring  $O(\log p)$  time.  $P_i$  continues to serve requests from other processors until receiving end-of-step signals from every processor, i.e., the end of the current step. At the end of each step,  $P_i$  receives  $|E_j|$  from each  $P_j$  by exchanging messages and it takes  $O(\log p)$  time.  $P_i$  updates  $q$  with the received  $|E_j|$ s in  $O(p)$  time. Then, in the next step,  $s$  number of edge switch operations are again distributed among  $p$  processors using multinomial distribution with the updated  $q$  and edge switch operations are performed. This process continues until  $t$  edge switch operations are performed in  $\lceil \frac{t}{s} \rceil$  steps.

#### 4.6. Properties of Parallel Edge Switch

In this section, we examine some stochastic properties of the parallel edge switch process and study how stochastically similar it is to the sequential edge switch process.

Recall that in the sequential edge switch process, one pair of edges is selected uniformly at random, and the edges are switched before selecting the next pair of edges. After completing the  $i$ -th edge switch operation, one or both of the two new edges generated by the  $i$ -th switch can be selected for the  $(i + 1)$ -th edge switch operation. In the parallel edge switch process, multiple pairs of edges are selected and switched simultaneously by different processors, and thus, the edges generated simultaneously by multiple processors cannot be selected for a simultaneous edge switch operation (restricting its choice). It raises the question of whether these two processes are stochastically equivalent or how close are they stochastically? We try to answer this question by studying the similarity of their effect, i.e., the resultant graphs generated by these two edge switch processes beginning with the same initial graph.

The *stochastic equivalence* of the sequential and parallel edge switch processes can be defined as follows. Let  $G_s^t$  and  $G_p^t$  be the resultant graphs after

performing  $t$  number of edge switch operations by the sequential and parallel edge switch processes, respectively, where both processes begin with the same initial graph  $G$ . We say the two processes are stochastically equivalent if  $\Pr\{G_s^t = G'\} = \Pr\{G_p^t = G'\}$  for all graphs  $G'$  with the same degree sequence as  $G$ .

Theoretical analysis of the above stochastic equivalence seems to be difficult. Experimental analysis can also be prohibitively time consuming. As the space of the graphs with a given degree sequence can be very large, estimating probabilities of generating  $G'$  by a reasonable number repetitions of the edge switch processes can be error prone.

Instead, we measure “similarity” of the two stochastic processes. We say the sequential and parallel processes are *similar* if they satisfy the following two conditions.

1. The distribution of the number of edges switched among different partitions (i.e., subsets of edges) is the same in both  $G_s^t$  and  $G_p^t$ , the resultant graphs of the sequential and parallel processes, respectively. This goal is achieved by the use of multinomial distribution as described before in Section 4.5.
2. At the end of the edge switch processes, the distribution of the number of edges across different sets of vertices is the same for both sequential and parallel processes. Let  $n_s(V_i, V_j)$  and  $n_p(V_i, V_j)$  be the number of cross edges between the sets of vertices  $V_i$  and  $V_j$  in the resultant graphs  $G_s^t$  and  $G_p^t$ , respectively. For any positive integer  $t$ , after switching  $t$  pairs of edges, the distribution of  $n_s(V_i, V_j)$  and  $n_p(V_i, V_j)$ , for all  $i, j$ , are same.

The resultant graphs,  $G_s^t$  and  $G_p^t$ , are divided into  $r$  partitions (i.e.,  $0 \leq i, j \leq r - 1$ ), with each partition containing an equal number of vertices having consecutive vertex labels. Note that the  $i$ -th partition  $V_i$  of  $G_s^t$  and  $G_p^t$  have the same set of vertices with vertex labels in  $\left[ \frac{i|V|}{r}, \frac{(i+1)|V|}{r} - 1 \right]$  (assuming  $n$  is a multiple of  $r$ ). The *edge difference*  $ED(G_s^t, G_p^t)$  across different sets of vertices between  $G_s^t$  and  $G_p^t$  is computed using equation (6). We define *error rate*

$ER(G_s^t, G_p^t)$  between  $G_s^t$  and  $G_p^t$  as shown in equation (7), where the maximum value of  $ED(G_s^t, G_p^t)$  can be  $2m$ . Due to randomness, some error rate can be observed even between two resultant graphs,  $G_{s1}^t$  and  $G_{s2}^t$ , generated by the sequential process in two different runs. If  $ER(G_s^t, G_p^t)$  is roughly equal to  $ER(G_{s1}^t, G_{s2}^t)$ , then the sequential and parallel processes are said to be *similar*. For a same pair of resultant graphs  $G_s^t$  and  $G_p^t$ , the value of  $ER(G_s^t, G_p^t)$  is different for different values of  $r$ . As a result, for a particular value of  $r$ , we are interested in how close  $ER(G_s^t, G_p^t)$  and  $ER(G_{s1}^t, G_{s2}^t)$  are to each other rather than the value of the error rate. The experimental results are explained in next section.

$$ED(G_s^t, G_p^t) = \sum_{i,j \geq i} |n_s(V_i, V_j) - n_p(V_i, V_j)| \quad (6)$$

$$ER(G_s^t, G_p^t) = \frac{ED(G_s^t, G_p^t)}{2m} \times 100\% \quad (7)$$

#### 4.7. Experimental Results

In this section, we present strong and weak scaling of our parallel algorithm, demonstrate the *similarity* of the sequential and parallel edge switch processes, and analyze the trade-offs between step-size, error rate and speedup.

**Experimental Setup.** We use a high performance computing cluster of 64 Intel Sandy Bridge compute nodes (Dell C6220). Each computing node consists of a dual-socket Intel Sandy Bridge E5-2670 2.60GHz 8-core processor (16 cores per node) and 64GB of 1600MHz DDR3 RAM. The computing nodes are interconnected by Qlogic QDR Infiniband interconnects. To implement our algorithm, we use MPICH2 implementation (version 1.9) of MPI.

**Datasets.** We use both real-world and artificial networks for the experiments. A summary of the networks is provided in Table 2. New York, Los Angeles, and Miami are synthetic, yet realistic social contact networks [21]. Each vertex represents a person in that city, and each edge represents any ‘physical’ contact between two persons within a 24 hour time period. Flickr is an image based online community network [22]. LiveJournal is a social network blogging

Table 2: Datasets used in the experiments.

Network	Type of network	Vertices	Edges	Avg. Degree
New York	Social Contact	20.38M	587.3M	57.63
Los Angeles	Social Contact	16.33M	479.4M	58.66
Miami	Social Contact	2.1M	52.7M	50.4
Flickr	Online Community	2.3M	22.8M	19.83
LiveJournal	Social	4.8M	42.8M	17.83
Small World	Random	4.8M	48M	20
Erdős-Rényi	Erdős-Rényi Random	4.8M	48M	20
PA-100M	Pref. Attachment	100M	1B	20
PA-1B	Pref. Attachment	1B	10B	20

site [22]. The small world graph is generated using the Watts-Strogatz small world graph model [23], Erdős-Rényi is generated using the Erdős-Rényi graph model [24], and PA is generated using the Preferential Attachment graph model [25].

**Strong Scaling.** Figure 4 showcases the strong scaling of the parallel algorithm of edge switch. The algorithm performs  $t$  edge switch operations to achieve a visit rate of  $x = 1$  using a step-size of  $t/100$ . We have experimented with eight different graphs, and achieved a maximum speedup of 85 using 1024 processors for the LiveJournal graph.

**Weak Scaling.** The weak scaling of our parallel algorithm is shown in Figure 5. In one experiment, we increase the graph size with the increase of processors, and use Preferential Attachment graphs with  $(p \times 0.1M)$  vertices and an average degree of 20. In another experiment, we use a fixed Preferential Attachment graph with  $102.4M$  vertices and  $1.024B$  edges. In both the experiments we use  $t = p \times 10M$  and step size =  $t/1000$ . Ideally, the parallel runtime should remain constant. However, in practice the communication increases with the increase of processors, leading to a higher runtime. Our algorithm shows good weak scaling as the runtime increases linearly in both the cases.

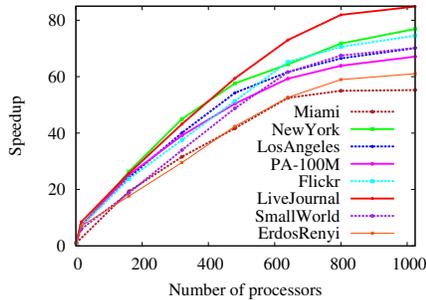


Figure 4: Strong scaling of our algorithm on eight different graphs using visit rate = 1 and step-size =  $t/100$ .

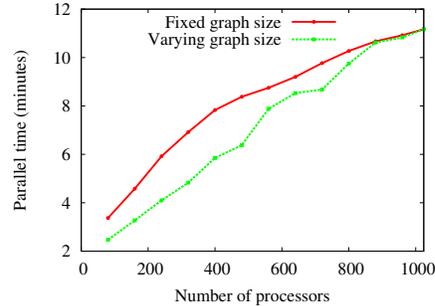


Figure 5: Weak scaling of our algorithm with fixed and varying size PA graphs. In one experiment, we use a fixed graph having  $102.4M$  vertices and  $1.024B$  edges while in the other experiment, we increase (or vary) graph size with the increase of processors. The varying graphs have  $(p \times 0.1M)$  vertices and an average degree of 20, where  $p$  is the number of processors. For both experiments, we use  $t = p \times 10M$  and step-size =  $t/1000$ .

**Similarity of the Outcomes of the Parallel and Sequential Algorithms and Determining Suitable Step-size.** We use visit rate  $x = 1$ , current calendar time as random seed,  $r = 20$  partitions,  $p = 1024$  processors, and average value of 10 experiments. Figure 6 shows that better strong scaling is achieved for a larger step-size on the Miami graph. For a particular step-size, error rate remains roughly constant with the increase of processors on the Miami graph, as shown in Figure 7. The effects of step-size on speedup and error rate for the Miami graph are shown in Figure 8 and Figure 9 respectively. Both the speedup and error rate increase with the increase of step-size.

While keeping the error rate to a minimum, we want to achieve as much speedup as possible. From Figure 9, we observe that with up to a  $2M$  step-size, the error rate between the resultant graphs generated by the sequential and parallel algorithms is roughly same as the error rate between the resultant graphs

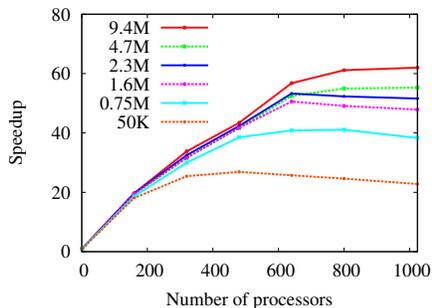


Figure 6: A comparison of strong scaling performance on Miami graph for different step-sizes: 9.4M, 4.7M, 2.3M, 1.6M, 0.75M, and 50K.

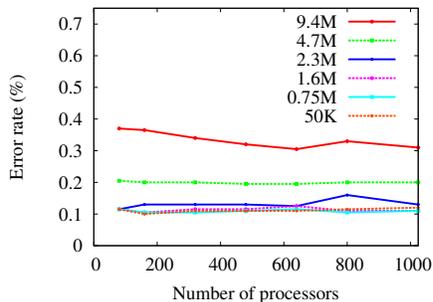


Figure 7: Error rate with increasing number of processors on Miami graph using different step-sizes: 9.4M, 4.7M, 2.3M, 1.6M, 0.75M, and 50K.

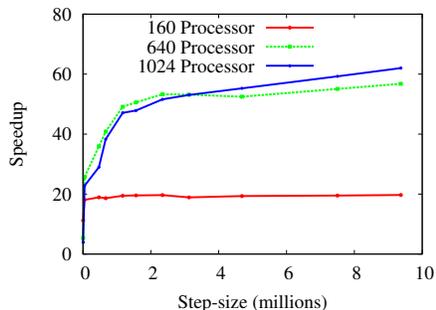


Figure 8: Speedup with increasing step-size on Miami graph using 160, 640 and 1024 processors.

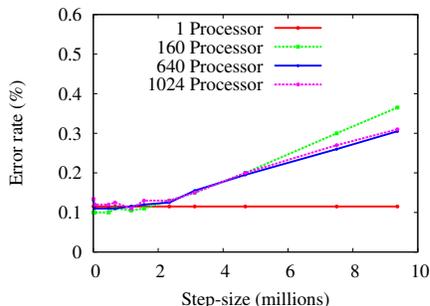


Figure 9: Error rate with increasing step-size on Miami graph using 1, 160, 640 and 1024 processors.

generated by two different execution of the sequential algorithm. Hence,  $2M$  can be a *suitable* step-size for the Miami graph, since the error rate is minimal, and a good speedup factor of 50 using 1024 processors is achieved at the same time. If we further increase the step-size, both the speedup and error rate increase. For example, using a step-size of  $9.4M$ , the error rate is a negligible  $0.4\%$ , however a higher speedup factor of 62 is achieved using 1024 processors. Figure 10 and 11 illustrate the effect of step-size on speedup and error rate, respectively, for

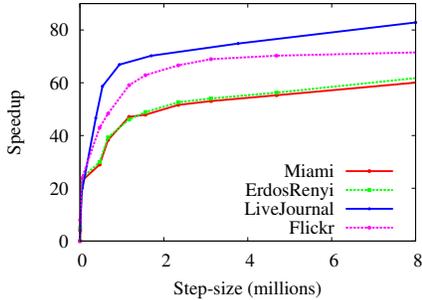


Figure 10: Speedup with increasing step-size for different graphs using 1024 processors.

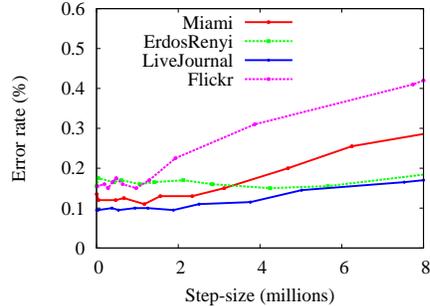


Figure 11: Error rate with increasing step-size for different graphs using 1024 processors.

different graphs. Suitable step-size may vary from graph to graph, depending on the graph size and type of the graph. For example, the error rate is roughly constant for different step-sizes on Erdős-Rényi and LiveJournal graphs, though it varies for Flickr and Miami graphs as shown in Figure 11. A *suitable* step-size for Flickr, Miami, LiveJournal and Erdős-Rényi graphs can be  $1.5M$ ,  $2M$ ,  $4M$  and  $8M$  respectively. In general, if we use a lower step-size, say  $2M$ , for any medium-sized graph (having more than  $20M$  edges), we expect to have a very small error rate along with a good speedup. The above experiments show that the sequential and the parallel edge switch processes are similar with a suitable step-size.

**How Network Properties Change with Edge Switching?** We also analyze how some network properties change with edge switch operations by the sequential and parallel algorithms. We use Miami, LiveJournal, and Flickr graphs with a step-size of  $2M$ , and vary the visit rate from 0.1 to 1. Figure 12 and 13 show that the average clustering coefficient and average shortest path distance of a graph change exactly the same way with edge switches by the sequential and parallel algorithms. Small variation in average shortest path distance is observed due to using approximate computation, since the exact computation is very time consuming.

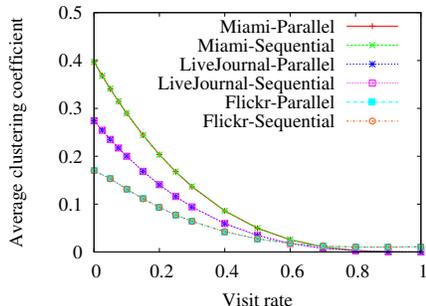


Figure 12: Average clustering coefficient changes similarly with edge switch operations by the sequential and parallel algorithms.

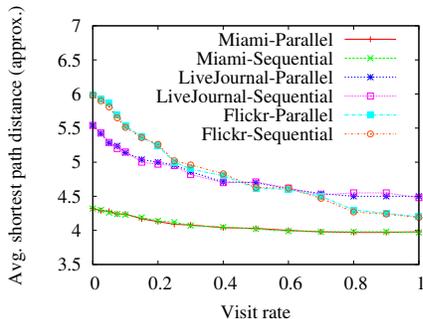


Figure 13: Average shortest path distance (approximate) changes similarly with edge switch operations by the sequential and parallel algorithms.

**Edge Switching in Large Networks.** Our parallel algorithm is able to perform more than  $115B$  edge switch operations on a Preferential Attachment graph with  $1B$  vertices and  $10B$  edges in less than 3 hours using 1024 processors. Due to the large size of the graph, the sequential algorithm is not even able to work on such a large graph.

## 5. Parallel Algorithms Using Hash-Based Partitioning

Partitioning schemes usually have good impact on the performance of parallel graph algorithms in terms of both runtime and memory. In the previous section, we presented a parallel algorithm for switching edges using a simple *consecutive partitioning (CP)* scheme. A good partitioning scheme for the parallel algorithm should have the following properties.

- It can efficiently partition a given network.
- Given a vertex  $v$ , the partition where  $v$  belongs to can be efficiently determined.
- The workload is uniformly distributed among the processors for different types of network.

The *workload* at a processor  $P_i$  is the number of edge switch operations  $P_i$  performs, which is proportional to the number of edges belonging to  $P_i$ . We use CP scheme because it reasonably fulfills all the criteria of a good partitioning scheme. In the CP scheme, a graph is partitioned such that each partition contains a subset of vertices having consecutive labels (or vertex-ids) and almost an equal number of edges; it is easy to determine which vertex belongs to which subset of consecutive vertices, therefore to which partition (or processor). It shows reasonably good performance as well. During the course of an edge switch process, the number of edges gradually change among the processors, which in turn makes the number of edge switches performed by processors skewed over time. As a result, in many cases, CP does not exhibit a well-balanced workload distribution among the processors. We discuss this phenomena later in this section. We also present a rigorous comparative study of several parallel algorithms using various partitioning schemes. Our experiments show that *hash-based partitioning (HP)* schemes demonstrate good performance in general as well as outperform the CP scheme in many cases. The rest of the section describes a few more parallel algorithms using several HP schemes and illustrates their performance. We further investigate the trade-offs between the HP and CP schemes with experimental results.

Among other options, one simple way to partition a given network is assigning vertices to partitions uniformly at random. This approach may assign almost an equal number of vertices to the partitions although the number of edges may vary among them. However, to determine which vertex belongs to which partition, each processor requires  $O(n)$  memory to map all vertex labels to partitions. Therefore assigning vertices arbitrarily among the processors may not be a good choice.

### 5.1. Hash-Based Partitioning

Another approach can be to use a hash-based partitioning scheme. A hash function can be a simple algebraic expression mapping vertex labels to partitions. Hash functions are deterministic in nature, and by using some simple hash

functions it can be very easy and efficient to determine which vertex belongs to which partition, thus obeying the first two criteria of a good partitioning scheme. Hash functions may assign different number of vertices and edges to partitions, although in many cases they may exhibit good workload distribution among the processors.

**Overview of the Parallel Algorithms using the HP Schemes.** The parallel algorithms using the HP schemes use the same data structure as before (Section 4.2). The graph is partitioned and distributed among the processors using the HP schemes. Once the graph is partitioned, the rest of the parallel algorithms remain the same for all schemes. That is,  $t$  edge switch operations are performed in a number of steps by  $p$  processors. At the beginning of a step,  $s$  edge switch operations are distributed among the processors by generating multinomial random variables in parallel. Each processor simultaneously performs edge switch operations as described in Section 4.4 and 4.5. At the end of every step, each processor updates the probability values of selecting edges from different partitions, which are then used in the next step of computation. All of the  $t$  edge switch operations are performed in such  $\lceil \frac{t}{s} \rceil$  steps.

**Partitioning the Network.** Given a hash function  $h$ , a vertex  $v$  is assigned to a partition belonging to processor  $P_i$  iff  $h(v) = i$ . The partition assigned to  $P_i$  contains a subset of vertices,  $V_i = \{v \in V | h(v) = i\}$  and a subset of edges,  $E_i = \{(u, v) \in E | u \in V_i, u < v\}$  such that  $\bigcup_i V_i = V$  and  $\bigcup_i E_i = E$ . Note that the partitions are disjoint, i.e., for  $i \neq j$ ,  $V_i \cap V_j = \phi$  and  $E_i \cap E_j = \phi$ .

A good hash function for the partitioning schemes should have the following properties.

- It is simple and efficient to determine which vertex belongs to which partition.
- Vertices are dispersed and well-distributed among the processors, i.e., all of the partitions are almost equal in size.

*Division hash function (HP-D), multiplication hash function (HP-M), and universal hashing (HP-U)* are few such hash functions and they are described below.

### 5.1.1. Division Hash Function

A simple hash function can be a division function (HP-D) [26]. This scheme uses the following function:

$$h(v) = v \bmod p \tag{8}$$

where  $p$  is the number of processors.

### 5.1.2. Multiplication Hash Function

Another simple hash function is a multiplication function (HP-M) [26]. The hash function is:

$$h(v) = \lfloor p(va - \lfloor va \rfloor) \rfloor \tag{9}$$

where  $a \in (0, 1)$  is a constant. The fractional part of  $va$  is extracted by  $va - \lfloor va \rfloor$  and is then multiplied by the number of processors  $p$  to determine the partition where  $v$  belongs to. Although this scheme works with any value of  $a \in (0, 1)$ , we use  $a = (\sqrt{5}-1)/2$  as suggested in [26] to obtain a reasonably good performance.

### 5.1.3. Universal Hashing

The division and multiplication hash functions are quite simple. However their workload distributions among the processors are dependent on the vertex labels of the input graph. If there is an adversary which knows the hash function being used in advance, the adversary can artificially manipulate the graph by assigning vertex labels in such a way that the workload distribution becomes skewed. For example, many high degree vertices can be assigned to a partition making the workload at the processor containing that partition significantly higher compared to other processors. To deal with such exploitation of hash functions by an adversary, universal hashing [26] can be a good choice. This scheme uses the following hash function:

$$h(v) = (((av + b) \bmod c) \bmod p) \tag{10}$$

where  $c$  is a large prime number such that all vertex labels are in the range  $[0, c - 1]$ ,  $a \in [1, c - 1]$  is a random integer, and  $b \in [0, c - 1]$  is another random

integer. Since  $a$  and  $b$  are selected randomly, this method arbitrarily selects a hash function from a large set of hash functions. As a result, there is no way for the adversary to know the exact hash function in advance, or exploit it to create a worse case scenario.

### 5.2. Experimental Results on Hash-Based Partitioning

In this section, we showcase the performance of the parallel algorithms using the HP schemes and demonstrate the trade-offs between the HP and CP schemes. The algorithms perform edge switch operations to achieve a visit rate of  $x = 1$  using a step-size of  $t/100$  for all experiments unless otherwise specified.

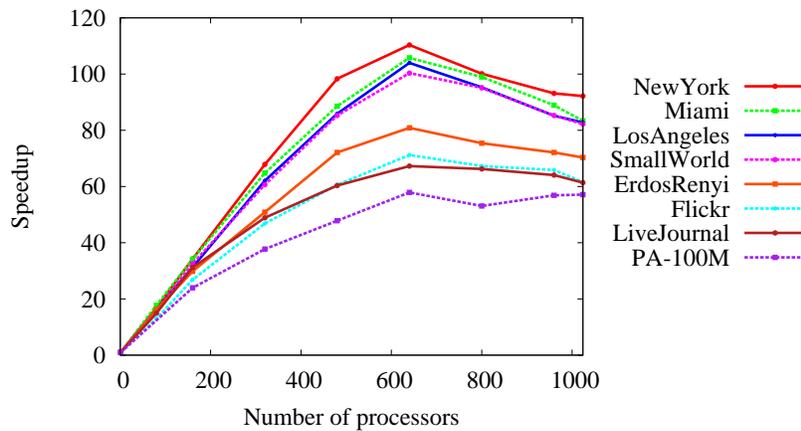


Figure 14: Strong scaling of the parallel algorithm using the HP-U partitioning scheme on eight different graphs.

**Strong Scaling.** Figure 14 illustrates strong scaling of the parallel algorithm using the HP-U partitioning scheme on eight different graphs. A good speedup of 110 is achieved using 640 processors on NewYork graph. Speedup varies for different graphs because of the types of graphs and difference in workload distribution among the processors. Speedup starts decreasing after some point with the increase of the number of processors indicating the domination of communication costs over computation costs.

A comparison of strong scaling performance of the parallel algorithms using

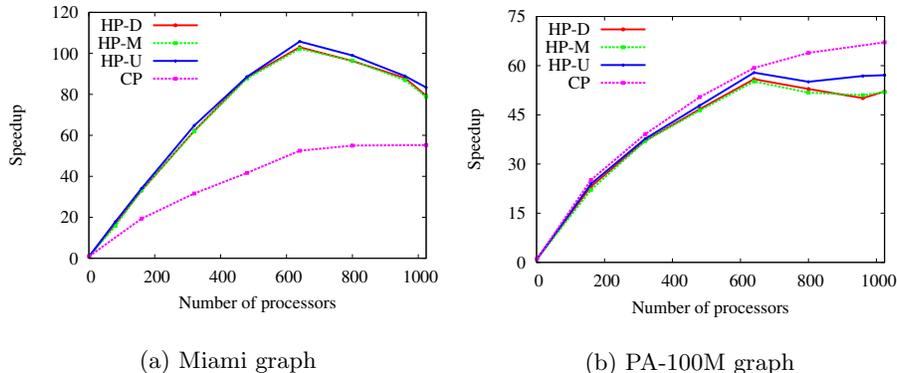


Figure 15: A comparison of strong scaling of the parallel algorithms using the HP-D, HP-M, HP-U and CP partitioning schemes for Miami and PA-100M graphs.

different schemes on Miami and PA-100M graph is demonstrated in Figure 15. The HP-D and HP-M schemes show good speedup as HP-U does. The HP-U scheme shows better strong scaling for Miami graph whereas CP outperforms the other schemes for PA-100M graph. To understand why speedup varies for different schemes and how good the schemes perform for different types of graphs, we further investigate workload distributions of different schemes on Miami and PA-100M graphs. We use  $p = 1024$  processors for rest of the experiments in this section.

**Load Balancing.** Figure 16 and 17 show distributions of vertices and edges (at the beginning of execution), respectively, among the processors in different schemes on Miami graph. The HP schemes assign roughly an equal number of vertices whereas the CP scheme initially assigns almost an equal number of edges among the processors. Due to the use of reduced adjacency lists, the number of vertices assigned to processors by CP scheme gradually increases with the increase of processor ranks despite having an equal number of edges among the processors. The number of edges initially assigned to all the processors by the HP schemes are very close to each other and the distribution can be considered as roughly load balanced although is not as perfect as that of the CP scheme.

All parallel algorithms start the edge switch process with almost an equal

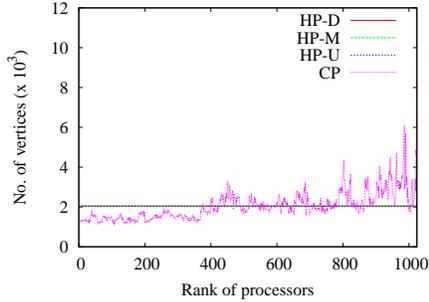


Figure 16: Distribution of vertices among the processors in HP-D, HP-M, HP-U and CP partitioning schemes for Miami graph.

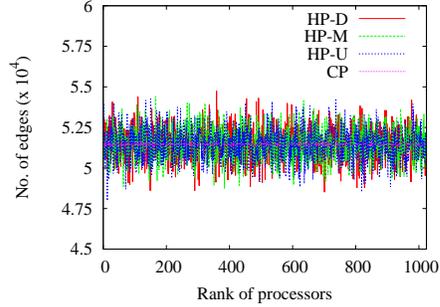


Figure 17: Distribution of edges (at the beginning of execution) among the processors in HP-D, HP-M, HP-U and CP partitioning schemes for Miami graph.

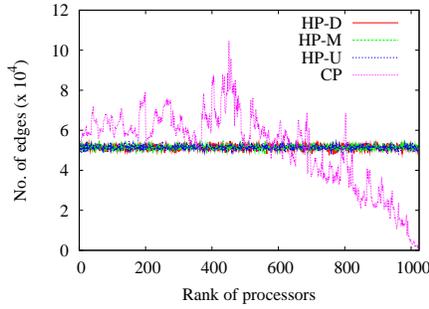


Figure 18: Distribution of edges (after completing execution) among the processors in HP-D, HP-M, HP-U and CP partitioning schemes for Miami graph.

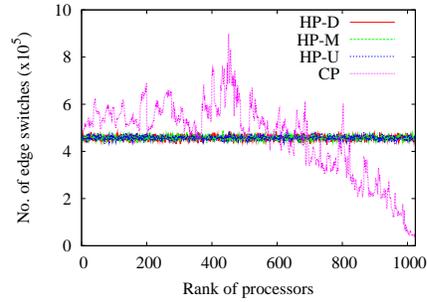


Figure 19: Distribution of workload (number of edge switch operations) among the processors in HP-D, HP-M, HP-U and CP partitioning schemes for Miami graph.

number of edges at each processor as shown in Figure 17. Recall that the number of edges gradually change among the processors with the progress of the edge switch process. As a result, at the completion of the edge switch process, the processors may end up with number of edges different than the numbers at the beginning of the process. Figure 18 shows the distribution of edges at the completion of an edge switch process using different schemes on Miami graph.

The CP scheme shows highly skewed distribution of edges compared to the HP schemes. The skewness exhibited in the CP scheme is a combined effect of the following reasons.

- A reduced adjacency list uses the ordering of vertex labels (from 0 to  $n-1$ ) to store an edge  $(u, v)$ :  $N(u)$  stores  $v$  if and only if  $u < v$ .
- The same ordering of vertex labels is used to assign a consecutive subset of vertices to a partition.

For example, let  $(u_1, v_1)$  be an edge belonging to the partition in the highest ranked processor  $P_{p-1}$ , participating in an edge switch operation with another edge  $(u_2, v_2)$  belonging to the partition in  $P_i$  ( $i < p-1$ ). There is a probability that both of the replacing edges (edge  $e_3$  and  $e_4$ , or  $e_5$  and  $e_6$  in Figure 3) can belong to  $N(u_2)$  and  $N(v_2)$ , which reside in some processors other than  $P_{p-1}$ , thus decreasing one edge from the partition in  $P_{p-1}$  and increasing one edge in the partition in  $P_j$  ( $j \neq p-1$ ). The occurrence of such scenario increases for graphs having a high clustering coefficient. Note that Miami is a synthetic yet realistic contact network with maximum, minimum, and average degree of 425, 1, and 50.4 respectively. It has a good clustering among the vertices that is gradually destroyed with progression of the edge switch process. For the Miami graph, most of the edges in the partition belonging to the highest ranked processor are replaced by edges with one end vertex belonging to some other partition, thus destroying the clustering among the vertices in the highest ranked processor as well as decreasing the number of edges in the partition substantially. As a result, some processors contain a higher number of edges compared to other processors at the end of the edge switch process. Since the number of edge switch operations performed at a processor  $P_i$  depends on the number of edges at  $P_i$ , the skewness in the number of edges among the processors with the course of the edge switch process results in an imbalanced workload distribution as shown in Figure 19 for the Miami graph.

In contrast, the HP schemes do not assign consecutive vertices to a partition. Thus a subset of vertices having dispersed vertex labels along with their reduced

adjacency lists belongs to a partition. As a result, the change in the number of edges among the partitions during the edge switch process is significantly less than that of the CP scheme for the Miami graph, leading to a better workload distribution in the HP schemes as shown in Figure 19. Hence, all of the HP schemes outperform the CP scheme for the Miami graph, which is illustrated in Figure 15. Among the three HP schemes, HP-U outperforms the others by a slight margin.

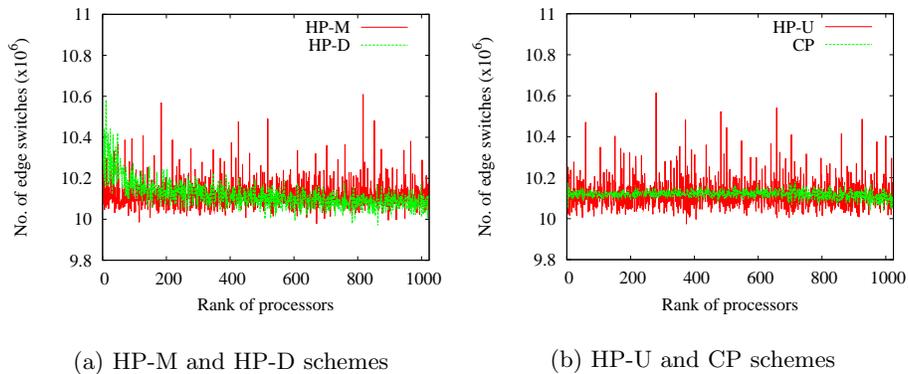


Figure 20: Distribution of workload (number of edge switch operations) among the processors in HP-D, HP-M, HP-U and CP partitioning schemes for PA-100M graph.

On the other hand, Figure 20 illustrates that the CP scheme exhibits better workload distribution for a Preferential Attachment graph having 100M vertices and 1B edges. PA graph has a very highly skewed degree distribution, i.e., it has few very high degree and many low degree vertices. The maximum, minimum, and average degree of PA-100M graph is 55225, 10, and 20 respectively. The CP scheme assigns a consecutive subset of vertices to partitions and uses the degrees of vertices to ensure that all the partitions have an equal number of edges; whereas the HP schemes assign vertices to partitions using only vertex labels; they neither use the degree of vertices nor consider the number of edges already assigned to a partition. As a result, the HP schemes assign several high degree vertices to some processors for the PA graph, thus making the initial edge distribution slightly more skewed compared to the CP scheme. Since

PA is a random graph having a very low clustering coefficient, the number of edges initially assigned to processors vary negligibly with the course of the edge switch process in the CP scheme. As a result, the CP scheme has an advantage of a better initial edge distribution, and thus demonstrates a better workload distribution and speedup compared to the HP schemes as shown in Figure 20 and 15 respectively.

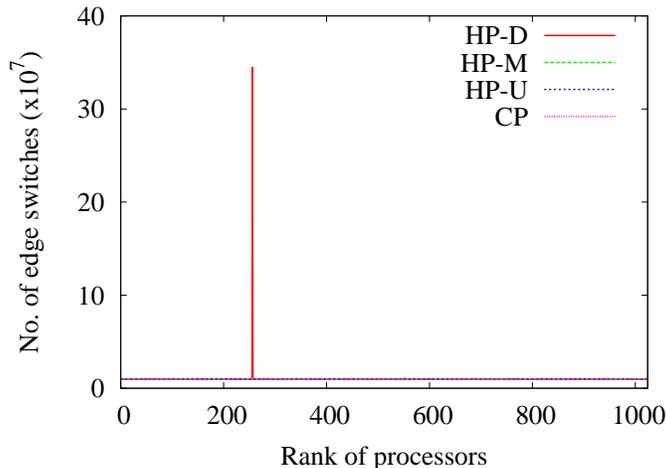


Figure 21: A worse case scenario of distribution of workload (number of edge switch operations) among the processors for the HP-D scheme on PA-100M graph.

**A Worse Case Scenario for the HP-D Scheme.** Unlike the CP scheme, where each partition contains almost an equal number of edges, one potential disadvantage of the HP schemes is that if there is an adversary aware of the exact hash function being used as the partitioning scheme, the adversary may generate a worse case scenario by artificially manipulating vertex labels of a graph. We simulate such a scenario for the HP-D scheme using 1024 processors. We intentionally reassign vertex labels of the PA-100M graph in such a way that all of the  $n/p$  highest degree vertices are assigned to a processor, say  $P_k$ . Thus  $P_k$  has a very high number of edges compared to other processors despite having an equal number of vertices among the processors. As a result,  $P_k$  performs substantially higher number of edge switch operations compared to

other processors as shown in Figure 21 (in this example,  $P_k$  is the processor with rank 256), whereas CP shows good performance by executing 28 times faster on the same graph as shown in Figure 22. An adversary can generate a similar worse case scenario for the HP-M scheme as well.

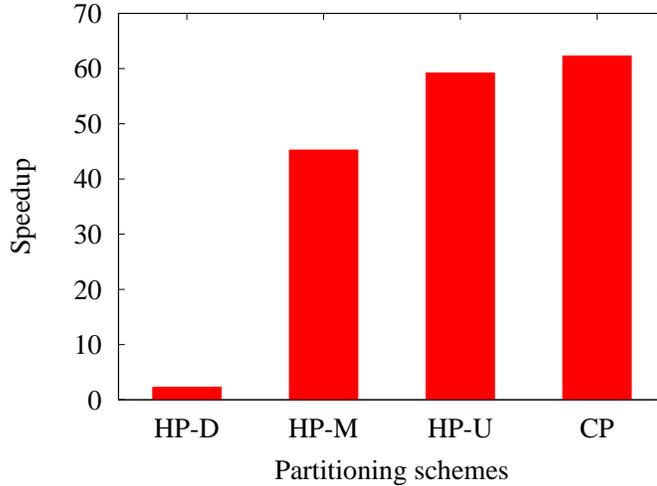


Figure 22: A comparison of speedup of a worse case scenario for the HP-D scheme with other schemes on a PA-100M graph using 1024 processors.

**Advantage of the HP-U Scheme.** The universal hashing randomly selects a hash function from a large set of hash functions. As a consequence, there is no way for an adversary to know in advance exactly which hash function will be used. Therefore, the HP-U scheme overcomes the drawbacks of the HP-D and HP-M schemes. In addition, HP-U demonstrates good speedup for all types of graphs and outperforms the other schemes in many cases. Although the CP scheme exhibits the best performance for PA-100M graph, speedup achieved by the HP-U scheme is very close to that of CP, justifying HP-U as a good choice in general.

**Weak Scaling.** Figure 23 shows weak scaling comparison of different schemes on PA graphs with the same experimental setup as before (Section 4.7). All of the schemes exhibit good weak scaling performance.

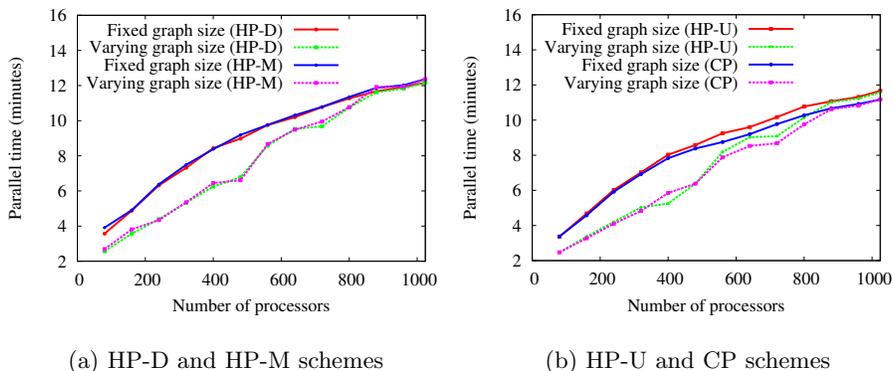


Figure 23: Weak scaling comparison of the parallel algorithms using the HP-D, HP-M, HP-U and CP partitioning schemes with fixed and varying size PA graphs. In one experiment, we use a fixed graph having  $102.4M$  vertices and  $1.024B$  edges while in the other experiment, we increase (or vary) graph size with the increase of processors. The varying graphs have  $(p \times 0.1M)$  vertices and an average degree of 20, where  $p$  is the number of processors. For both experiments, we use  $t = p \times 10M$  and step-size =  $t/1000$ .

**Similarity of the Outcomes of the Parallel and Sequential Algorithms and Determining Suitable Step-size.** To study the similarity of the outcomes of the algorithms, we use the same experimental setup as before. That is, we use visit rate  $x = 1$ ,  $p = 1024$  processors for the parallel algorithms, current calendar time as random seed,  $r = 20$  partitions, and average values of 10 experiments to measure the error rate. Table 3 shows the error rate comparison of the outcomes of the parallel algorithms using different schemes, with that

Table 3: Error rate comparison of the outcomes of the parallel algorithms using HP-D, HP-M, HP-U and CP partitioning schemes with that of the sequential algorithm for different graphs. We use average values of 10 experiments.

Networks	Error rate (%) for different schemes							
	Sequential	Parallel						
		Using 1 step			Using 100 steps			
		HP-D	HP-M	HP-U	HP-D	HP-M	HP-U	CP
Miami	0.117	0.118	0.123	0.117	0.111	0.127	0.123	0.164
SmallWorld	0.112	0.100	0.112	0.119	0.106	0.118	0.109	0.115
LiveJournal	0.116	0.117	0.118	0.117	0.116	0.116	0.116	0.115

of the sequential one suggesting that even for performing edge switch operations in one step, the outcomes of the parallel algorithms using the HP schemes are similar to that of the sequential algorithm with a negligible error rate deviation. Since the HP schemes assign vertices dispersedly among the partitions, the number of edges initially belonging to the partitions change negligibly with edge switch operations compared to that of the CP scheme. Hence the HP schemes can perform edge switch operations in only one step with reasonable accuracy, which consequently makes computation faster. As a result, the parallel algorithms using the HP schemes no longer need a suitable step-size. In contrast, finding a suitable step-size is important for the CP scheme to obtain a close approximation of the outcome of the sequential algorithm.

Our experiments show that the average clustering coefficient and average shortest path distance of a graph change with edge switch operations by the parallel algorithms exactly the same way for all the schemes as they change with edge switches by the sequential algorithm. The outcomes are similar to the results shown in Figure 12 and 13.

**Synopsis of the Experimental Results.** All of the partitioning schemes demonstrate reasonably good performance. Below is a summary of the results.

- The hash-based partitioning schemes exhibit better performance for many graphs (e.g., Miami) because of a well-balanced workload distribution.
- The HP schemes can perform edge switch operations in only one step with reasonable accuracy, thus eliminating the need for performing edge switch operations in a number of steps.
- There is a possibility of a worst case scenario arising for the HP-D and HP-M schemes that may slow down the algorithms significantly. The HP-U scheme overcomes this drawback by randomly choosing a hash function from a large set of hash functions, and illustrates good on-average performance as well as outperforming the rest of the schemes in many cases.
- The CP scheme shows good performance with some computation overhead

by performing edge switch operations with a suitable step-size for all types of graphs, and in some cases (e.g., PA-100M) outperforms the HP schemes.

- The CP scheme is not vulnerable to adversaries for generating worse case scenarios as the HP-D and HP-M schemes do.

## 6. Parallel Algorithm for Computing Binomial and Multinomial Distribution

In this section we present a parallel algorithm for computing multinomial distribution of very large numbers. First we briefly review the current state-of-the-art sequential algorithm.

### 6.1. Sequential Algorithm for Computing Multinomial Distribution

One simple approach for computing multinomial random variables is to perform  $N$  independent trials, where the outcome of each trial can be  $0, 1, \dots, \ell - 1$  with probability  $q_0, q_1, \dots, q_{\ell-1}$ , respectively. This algorithm takes at least  $\Omega(N \log \ell)$  time. An efficient state-of-the-art algorithm is the *conditional distributed method* [20], which runs in  $O(N)$  time. This method generates multinomial random variables  $\langle X_0, X_1, \dots, X_{\ell-1} \rangle$  by iteratively generating  $\ell$  binomial random variables:

$$X_i \sim \mathcal{B} \left( N - \sum_{j=0}^{i-1} X_j, \frac{q_i}{1 - \sum_{j=0}^{i-1} q_j} \right) \quad (11)$$

Inverse transformation method (BINV) [27] is the best known algorithm for computing binomial random variables. To generate a binomial random variable  $X$  with parameters  $N$  and  $q$ , it takes  $O(X)$  time. Note that the expected value of  $X$  is  $Nq$ .

The algorithms for the inverse transformation method (BINV) [27] to generate binomial random variables and for the conditional distributed method [20] to generate multinomial random variables are shown in Algorithm 3 and 4, respectively. For additional details, see [27] and [20].

---

**Algorithm 3** BINOMIAL( $N, q$ )

---

```
1: if  $q = 1$  then return  $N$ 
2:  $i \leftarrow 0$    $\{i$  is the binomial random variable $\}$ 
3: Generate  $u \sim U(0, 1)$  uniformly at random
4:  $Q \leftarrow (1 - q)^N$ ,  $S \leftarrow Q$ 
5: while  $S < u$  do
6:    $i \leftarrow i + 1$ 
7:    $Q \leftarrow Q \binom{N-i+1}{i} \left(\frac{q}{1-q}\right)$ 
8:    $S \leftarrow S + Q$ 
9: return  $i$ 
```

---

---

**Algorithm 4** MULTINOMIAL( $N, q_0, q_1, \dots, q_{\ell-1}$ )

---

```
1:  $X_s \leftarrow 0, Q_s \leftarrow 0$ 
2: for  $i = 0$  to  $\ell - 1$  do
3:   if  $Q_s < 1$  then
4:      $X_i \leftarrow \text{BINOMIAL}\left(N - X_s, \frac{q_i}{1 - Q_s}\right)$ 
5:      $X_s \leftarrow X_s + X_i$ 
6:      $Q_s \leftarrow Q_s + q_i$ 
7:   else  $X_i \leftarrow 0$ 
8: return  $\langle X_0, X_1, \dots, X_{\ell-1} \rangle$ 
```

---

The conditional distributed method shown in Algorithm 4 runs in  $\sum_{i=0}^{\ell-1} O(X_i) = O(N)$  time. In the next section, we present an efficient parallelization of Algorithm 4.

### 6.2. Parallel Algorithm for Computing Multinomial Distribution

Based on the conditional distributed method shown in Algorithm 4, we propose a parallel algorithm for computing multinomial distribution  $X \sim \mathcal{M}(N, q)$ , where  $q$  denotes probability vector  $\langle q_0, q_1, \dots, q_{\ell-1} \rangle$ . One tempting approach to parallelize the conditional distributed method is to distribute the generation of  $X_i$ ,  $0 \leq i < \ell$  (Line 4 of Algorithm 4) among the processors. However, a difficulty arises from the sequential nature of computing  $X_i$ s due to the dependencies of  $X_i$  on  $X_{i-1}$  for all  $i > 0$ . We overcome this difficulty by exploiting some properties of binomial and multinomial random variables, as described below.

Let  $N_i$ , for  $0 \leq i < k$ , be some integers such that  $N = \sum_{i=0}^{k-1} N_i$ . If  $X_i \sim \mathcal{B}(N_i, q)$ , then

$$X = \sum_{i=0}^{k-1} X_i \sim \mathcal{B}\left(\sum_{i=0}^{k-1} N_i, q\right) = \mathcal{B}(N, q) \quad (12)$$

The above property of the binomial random variables leads to the following property of the multinomial random variables. If

$$\langle X_{0,i}, X_{1,i}, \dots, X_{\ell-1,i} \rangle \sim \mathcal{M}(N_i, q_0, q_1, \dots, q_{\ell-1})$$

for  $0 \leq i < k$ , then

$$\langle X_0, X_1, \dots, X_{\ell-1} \rangle \sim \mathcal{M}(N, q_0, q_1, \dots, q_{\ell-1}) \quad (13)$$

where  $X_j = \sum_{i=0}^{k-1} X_{j,i}$  for  $0 \leq j < \ell$  and  $N = \sum_{i=0}^{k-1} N_i$ .

Now we describe the parallel algorithm for computing multinomial distribution, which uses the above property. First, we explain the case of  $p = \ell$ . Our algorithm divides the number of trials  $N$  into  $p$  almost equal small number of trials  $N_i$ , and assign  $N_i$  to  $P_i$ . Then each processor  $P_i$  computes the multinomial distribution of  $N_i$  using the same probability vector  $q$ . At the end, the

results of all the processors are aggregated. The pseudocode is given in Algorithm 5, where processor  $P_i$  holds the multinomial random variable  $X_i$  at the end of computation.

---

**Algorithm 5** PARALLEL MULTINOMIAL( $N, q_0, \dots, q_{\ell-1}$ )

---

- 1: Each processor  $P_i$  executes the following in parallel:
  - 2: **if**  $i < N \% p$  **then**  $N_i \leftarrow \lfloor \frac{N}{p} \rfloor + 1$
  - 3: **else**  $N_i \leftarrow \lfloor \frac{N}{p} \rfloor$
  - 4:  $\langle X_{0,i}, X_{1,i}, \dots, X_{\ell-1,i} \rangle \sim \mathcal{M}(N_i, q_0, q_1, \dots, q_{\ell-1})$
  - 5: Send  $X_{j,i}$  to processor  $P_j$
  - 6: Upon receiving  $X_{i,k}$  from every processor  $P_k$ :
  - 7:  $X_i \leftarrow \sum_{k=0}^{p-1} X_{i,k}$
- 

For  $p \neq \ell$ , the algorithm is the same up to the multinomial distribution computation of  $N_i$  at  $P_i$ , i.e., lines 1-4 of Algorithm 5. The only difference is how the generated multinomial random variables will be stored among the processors. The variables can be stored in many ways, e.g., all the  $X_i$ s can be gathered to the root processor  $P_0$ , or they ( $X_i$ s) can be distributed among the processors in a round robin fashion, i.e., assigning  $X_i$  to processor  $P_{(i \% p)}$ , etc.  $X_i$  is always computed by summing up all the  $X_{i,k}$ s ( $0 \leq k < p$ ), after receiving them from all processors.

The parallel computation is almost perfectly load balanced among the processors since each processor computes multinomial distribution of  $N/p$  independently, taking  $O\left(\frac{N}{p}\right)$  time. The communication cost at the end takes  $O(\ell \log p)$  time. Hence, the time complexity of this algorithm is  $O\left(\frac{N}{p} + \ell \log p\right)$ . The algorithm is almost perfectly parallelized because the number of processors,  $p$  (which is in the range of hundreds or at most thousands), and the number of outcomes  $\ell$ , are significantly smaller than the number of trials  $N$  (which is in the range of billions), in a general case. Algorithm 5 computes binomial distribution for  $\ell = 2$ .

During binomial random variable generation, the computation of  $(1 - q)^N$  (Line 4 of Algorithm 3) results in underflow occurrence for large values of  $N$ , e.g., billions. Using *long double* data type cannot solve this underflow occurrence for large  $N$ . In addition, some round off errors may appear. We deal with these difficulties by using the property of the binomial distribution again, i.e., we divide  $N$  into small  $N_i$ s such that  $\sum_i N_i = N$ , compute  $X$  using equation (12). The upper threshold value of  $N_i$  is set such that no underflow occurs, that is,

$$(1 - q)^{N_i} \geq z \tag{14}$$

$$N_i \leq \frac{-\log z}{\log(1 - q)} \leq \frac{-\log z}{2q} \tag{15}$$

where  $z$  is the smallest positive real number that can be represented by the data type (e.g., float, double) used and  $q < 1$ .

### 6.3. Performance Analysis of the Parallel Algorithm

In this section, the performance of the parallel algorithm for multinomial distribution is demonstrated by strong scaling and weak scaling.

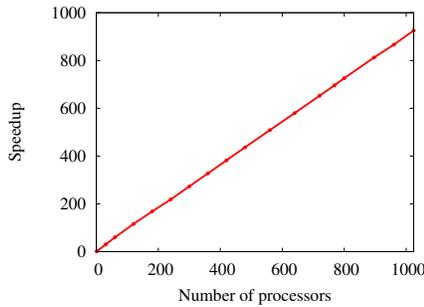


Figure 24: Strong scaling of the parallel algorithm of multinomial distribution using  $N = 10000B$ ,  $\ell = 20$  and  $q_i = 1/\ell$ .

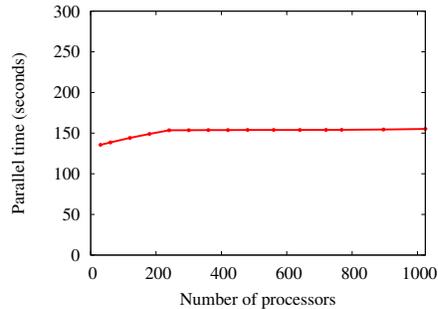


Figure 25: Weak scaling of the parallel algorithm of multinomial distribution using  $N = p \times 20B$ ,  $\ell = p$  and  $q_i = 1/\ell$ .

**Strong Scaling.** The strong scaling of the parallel algorithm is illustrated in Figure 24. We keep the problem size fixed ( $N = 10000B$ ,  $\ell = 20$  and  $q_i = 1/\ell$ ), and achieve a speedup of 925 using 1024 processors. The speedup

increases almost linearly with the increase of processors. The parallel algorithm can compute a multinomial distribution of  $10000B$  in 71 seconds using 1024 processors.

**Weak Scaling.** Figure 25 shows the weak scaling of our parallel algorithm. We use  $\ell = p$  (i.e., total number of processors),  $N = p \times 20B$  (i.e.,  $20B$  per processor), and equal probability values,  $q_i = 1/\ell$ . The parallel run time is almost constant indicating a very good weak scaling.

## 7. Conclusion

We presented several parallel algorithms for switching edges in massive networks. They can be used in studying various properties of large dynamic networks as well as in generating massive scale random graphs. The algorithms scale well to a large number of processors and exhibit good speedup. We also presented the trade-offs of the algorithms. In addition, we developed a parallel algorithm for generating multinomial random variables that is almost perfectly parallelized. This algorithm can be of independent interest and prove useful in parallelizing many other stochastic processes. We believe that the parallel algorithms will contribute significantly when dealing with big data, one of the most challenging problems in today's research world.

## Acknowledgment

We thank our external collaborators, members of the Network Dynamics and Simulation Science Laboratory (NDSSL), and anonymous reviewers for their suggestions and comments. We are grateful to Anil Vullikanti for interesting discussions and helpful comments on a draft of this paper. We also sincerely thank Maureen Lawrence-Kuether and Jim Walke for proof-reading this paper.

## References

- [1] H. Bhuiyan, J. Chen, M. Khan, M. Marathe, Fast parallel algorithms for edge-switching to achieve a target visit rate in heterogeneous graphs, in:

- Proceedings of the 43rd International Conference on Parallel Processing (ICPP), IEEE, 2014, pp. 60–69.
- [2] C. Cooper, M. Dyer, C. Greenhill, Sampling regular graphs and a peer-to-peer network, *Combinatorics, Probability and Computing* 16 (4) (2007) 557–593.
- [3] T. Feder, A. Guetz, M. Mihail, A. Saberi, A local switch markov chain on given degree graphs with application in connectivity of peer-to-peer networks, in: *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006, pp. 69–76.
- [4] C. Gkantsidis, M. Mihail, E. Zegura, The markov chain simulation method for generating connected power law random graphs, in: *Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments (ALENEX)*, Vol. 111, SIAM, 2003, pp. 16–25.
- [5] M. Jerrum, A. Sinclair, Fast uniform generation of regular graphs, *Theoretical Computer Science* 73 (1) (1990) 91–100.
- [6] R. Kannan, P. Tetali, S. Vempala, Simple markov-chain algorithms for generating bipartite graphs and tournaments, *Random Structures and Algorithms* 14 (4) (1999) 293–308.
- [7] J. Ray, A. Pinar, C. Seshadhri, Are we there yet? When to stop a markov chain while generating random graphs, in: *Proceedings of the 9th Workshop on Algorithms and Models for the Web Graph (WAW)*, Springer, 2012, pp. 153–164.
- [8] I. Stanton, A. Pinar, Constructing and sampling graphs with a prescribed joint degree distribution, *Journal of Experimental Algorithmics (JEA)* 17 (3) (2012) 3.5:3.1–3.5:3.25.
- [9] A. Stauffer, V. Barbosa, A study of the edge-switching markov-chain method for the generation of random graphs, *Tech. Rep. cs.DM/0512.105* (2005).

- [10] L. Tabourier, C. Roth, J. Cointet, Generating constrained random graphs using multiple edge switches, *Journal of Experimental Algorithmics (JEA)* 16 (1) (2011) 1.7:1.1–1.7:1.15.
- [11] M. Newman, The structure and function of complex networks, *SIAM Review* 45 (2) (2003) 167–256.
- [12] J. Blitzstein, P. Diaconis, A sequential importance sampling algorithm for generating random graphs with prescribed degrees, *Internet Mathematics* 6 (4) (2011) 489–522.
- [13] N. Wormald, *Models of random regular graphs*, London Mathematical Society Lecture Note Series (1999) 239–298.
- [14] A. Steger, N. Wormald, Generating random regular graphs quickly, *Combinatorics, Probability and Computing* 8 (04) (1999) 377–396.
- [15] J. Kim, V. Vu, Sandwiching random graphs: universality between random graph models, *Advances in Mathematics* 188 (2) (2004) 444–469.
- [16] S. Hakimi, On realizability of a set of integers as degrees of the vertices of a linear graph, *Journal of the Society for Industrial and Applied Mathematics* 10 (3) (1962) 496–506.
- [17] S. Eubank, A. Vullikanti, M. Khan, M. Marathe, C. Barrett, Beyond degree distributions: Local to global structure of social contact graphs, in: *Proceedings of the Third International Conference on Social Computing, Behavioral Modeling, and Prediction (SBP)*, 2010, p. 1.
- [18] A. Hagberg, P. Swart, D. Schult, Exploring network structure, dynamics, and function using NetworkX, in: *Proceedings of the 7th Python in Science Conference (SciPy)*, 2008, pp. 11–15.
- [19] I. Adler, S. Oren, S. Ross, The coupon-collector’s problem revisited, *Journal of Applied Probability* 40 (2) (2003) 513–518.

- [20] C. Davis, The computer generation of multinomial random variates, *Computational Statistics & Data Analytics* 16 (2) (1993) 205–217.
- [21] C. Barrett, R. Beckman, M. Khan, V. Kumar, M. Marathe, P. Stretz, T. Dutta, B. Lewis, Generation and analysis of large synthetic social contact networks, in: *Proceedings of the 2009 Winter Simulation Conference (WSC)*, 2009, pp. 1003–1014.
- [22] J. Leskovec, A. Krevl, SNAP Datasets: Stanford large network dataset collection, <http://snap.stanford.edu/data> (Jun. 2014).
- [23] D. Watts, S. Strogatz, Collective dynamics of ‘small-world’ networks, *Nature* 393 (6684) (1998) 440–442.
- [24] B. Bollobás, *Random graphs*, Springer, 1998.
- [25] A. Barabási, R. Albert, Emergence of scaling in random networks, *Science* 286 (5439) (1999) 509–512.
- [26] T. Cormen, *Introduction to algorithms*, MIT press, 2009.
- [27] V. Kachitvichyanukul, B. Schmeiser, Binomial random variate generation, *Communications of the ACM* 31 (2) (1988) 216–222.