

PASGAL: Parallel And Scalable Graph Algorithm Library

Xiaojun Dong
UC Riverside
xdong038@ucr.edu

Yan Gu
UC Riverside
ygu@cs.ucr.edu

Yihan Sun
UC Riverside
yihans@cs.ucr.edu

Letong Wang
UC Riverside
lwang323@ucr.edu

Abstract

In this paper, we introduce PASGAL (Parallel And Scalable Graph Algorithm Library), a parallel graph library that scales to a variety of graph types, many processors, and large graph sizes. One special focus of PASGAL is the efficiency on *large-diameter graphs*, which is a common challenge for many existing parallel graph processing systems: many existing graph processing systems can be even slower than the standard sequential algorithm on large-diameter graphs due to the lack of parallelism. Such performance degeneration is caused by the high overhead in scheduling and synchronizing threads when traversing the graph in the breadth-first order.

The core technique in PASGAL to achieve high parallelism is a technique called *vertical granularity control (VGC)* to hide synchronization overhead, as well as careful redesign of parallel graph algorithms and data structures. In our experiments, we compare PASGAL with state-of-the-art parallel implementations on BFS, SCC, BCC, and SSSP. PASGAL achieves competitive performance on small-diameter graphs compared to the parallel baselines, and is significantly faster on large-diameter graphs.

CCS Concepts

• **Theory of computation** → **Graph algorithms analysis**; **Parallel algorithms**; **Shared memory algorithms**.

Keywords

Parallel Algorithms, Graph Algorithms, Graph Processing

1 Introduction

Graphs are effective representations of real-world objects and their relationships. Processing and analyzing graphs efficiently have become increasingly important. Given the large size of today’s real-world graphs, it is imperative to consider parallelism in graph processing. The increasing number of cores and memory size allows a single machine to easily process graphs with billions of vertices in a few seconds, even for reasonably complicated tasks. As a result, a huge number of in-memory graph processing algorithms and systems have been developed (e.g. [5, 9]).

Despite the hardware advances, the increasing number of cores does not provide “free” performance improvement. We observed that many existing parallel systems suffer from scalability issues, both to more cores and to larger/more diverse graphs, even in fundamental tasks such as breadth-first search (BFS), strongly connected components (SCC), biconnected components (BCC), and single source shortest paths (SSSP). We show an example of SCC algorithms in Fig. 1. Tested on a 96-core machine, existing systems [9, 20] scale well with fewer than 16 threads and/or on the well-studied power-law graphs with small diameters. Indeed, many of them focus on optimizing the performance of low-diameter graphs such as social networks. However, their performance can stop increasing (or even drop) with more threads, especially on large-diameter graphs. On many graphs, they can perform worse than the sequential Tarjan’s algorithm [21]. Similar issues on other graph problems can be observed in Fig. 2.




Figure 1: Speedup on #processors for strongly connected component (SCC) algorithms over Tarjan’s algorithm [21] (sequential, always 1). The tested algorithms include PASGAL (this paper), GBBS [9], and Multistep [20]. We present four graphs of different types. The top two graphs have small diameters and the bottom two have larger diameters.

One major reason for such performance degeneration is the high overhead of managing and synchronizing threads. While enabling the potential of better parallelism, more available cores also bring up great challenges and overhead. Namely, *parallelism comes at a cost*. This is more pronounced when using BFS-like primitives on large-diameter graphs: when traversing the graph in BFS order, the number of rounds (and thus the cost of scheduling and synchronizing threads between them) is proportional to the diameter of the graph. As a result, when the diameter is large, the synchronization overhead can be more expensive than the computation.

We propose a new open-source library **PASGAL: the Parallel And Scalable Graph Algorithm Library**, that implements a list of graph algorithms that are scalable to diverse graph types, many processors, and large graphs. PASGAL focuses on several problems where existing parallel solutions suffer severely from high synchronization costs, such as BFS, SCC, BCC, and SSSP. We plan to include more algorithms in the future.

To overcome the scalability issues, the key technique in PASGAL is called *vertical granularity control (VGC)* proposed in our recent paper [24] to hide scheduling overhead. Accordingly, we need to redesign algorithms and data structures to facilitate VGC, as well as to synergistically optimize work, span and/or space usage. In Sec. 2, we introduce our techniques in more details.

With VGC and other techniques in Sec. 2, PASGAL achieves high performance on a variety of graphs, especially large-diameter graphs. For the aforementioned SCC problem, PASGAL exhibits good scalability (see Fig. 1), and is faster than all previous parallel algorithms on all tested graphs (see Fig. 2). Overall, compared to the baselines, PASGAL is always competitive on small-diameter graphs, and is almost always the fastest on large diameter graphs. We discuss experimental results in Sec. 3. Our code is publicly available [10]. Full experimental results and more references are given in the Appendix.

Preliminaries. Given a graph $G = (V, E)$, we denote $n = |V|$ and $m = |E|$. We use D to denote the diameter of G .

Algorithm 1: Parallel Frontier-based Graph Algorithms**Input:** A directed graph $G = (V, E)$ and an initial frontier $\mathcal{F}_0$

```

1  $i \leftarrow 0$ 
2 while  $\mathcal{F}_i \neq \emptyset$  do
3   parallel_for_each  $v \in \mathcal{F}_i$  do
4     Visit all neighbors of  $v$ , put a subset of them to  $\mathcal{F}_{i+1}$ .
5    $i \leftarrow i + 1$ 

```

Most algorithms in PASGAL are *frontier-based* (see Alg. 1). At a high level, the algorithm maintains a *frontier*, which is a subset of vertices to be explored in each round. In round i , the algorithm *processes* (visits their neighbors) the current frontier $\mathcal{F}_i$ in parallel, and puts a subset of their neighbors to the next frontier $\mathcal{F}_{i+1}$, determined by a certain condition. For example, in BFS, a vertex u will add its neighbor v to the next frontier iff. v has not been added to frontiers before, and u is the first to add v (based on some linearization order) in round i . The algorithm requires $\Omega(D)$ rounds.

As mentioned, one key challenge of using parallel BFS or similar approaches is the large cost to create and synchronize threads between rounds, which is especially costly for large-diameter graphs (more rounds needed). In this paper, we will show how PASGAL reduces the scheduling overhead to achieve better parallelism.

2 Algorithms

We now introduce the algorithms in PASGAL. For page limit, we only elaborate on SCC and briefly overview the others.

2.1 Parallel SCC

Most existing parallel SCC algorithms are based on *reachability search*, which finds all vertices u that are reachable from a given vertex v . In most existing implementations, reachability searches are performed by BFS from v , which requires $O(D)$ rounds to finish. This directly implies several (interrelated) challenges on large-diameter graphs. First, this incurs many rounds of distributing and synchronizing threads with high overhead. Second, many real-world large-diameter graphs (e.g., road networks) are sparse with small average degrees. As a result, every parallel task (processing one vertex in the frontier) is small, and the cost of scheduling the thread may dominate the actual computation. Finally, because of sparsity, each frontier size is also likely small, which makes the algorithm unable to utilize all threads in the hardware.

Algorithm Redesign. To resolve this challenge, PASGAL uses a recent SCC algorithm [24]. The idea is to observe that a reachability search does not require a strong BFS order. Therefore, one can relax the BFS order and visit vertices in an arbitrary order. In this way, the algorithm employs a technique called *vertical granularity control* to hide scheduling overhead, as introduced below.

Vertical Granularity Control. *Granularity control* (a.k.a. coarsening) is widely used in parallel programming, which also aims to avoid the overhead caused by generating unnecessary parallel tasks. For computations with sufficient parallelism, e.g., a parallel for-loop or divide-and-conquer algorithm of size $n \gg P$ where P is the number of processors, a common practice is to stop recursively creating parallel tasks at a certain subproblem size and switch to a sequential execution to hide the scheduling overhead.

Inspired by this, the idea of VGC is also to increase each task size to hide the scheduling overhead. For reachability searches, we simply perform a *local search* to visit at least τ vertices, possibly in multiple hops. While globally the vertices are not visited in the BFS order, this does not affect the correctness of reachability. Note that here τ is equivalent to the base-case task size of granularity control,

	Road				Social				
	AF	NA	AS	EU	LJ	FB	OK	TW	FS
n	33.5M	87.0M	95.7M	131M	4.85M	59.2M	3.07M	41.7M	65.6M
m'	44.8M	113M	123M	169M	69.0M	N/A	N/A	1.47B	3.61B
m	88.9M	220M	244M	333M	85.7M	185M	234M	2.41B	3.61B
D'	8276	9337	16660	11814	22	N/A	N/A	18	N/A
D	3948	4835	8794	4410	19	22	9	22	37

	k -NN				Web				
	CH5	GL5	GL10	COS5	WK	SD	CW	HL14	HL12
n	4.21M	24.9M	24.9M	321M	6.63M	89.2M	978M	1.72B	3.56B
m'	21.0M	124M	249M	1.61B	165M	2.04B	42.4B	64.4B	129B
m	29.7M	157M	310M	1.96B	300M	3.88B	74.7B	124B	226B
D'	5683	17268	13982	1390	62	145	506	800	5279
D	14479	21601	9053	1180	9	35	254	366	650

	Synthetic				Underline: undirected graphs				
	REC	SREC	TRCE	BBL	m' : #edges in directed graphs				
n	100M	100M	16.0M	21.2M	m : #edges in undirected or				
m'	297M	204M	N/A	N/A	symmetrized graph				
m	400M	336M	48.0M	63.6M	D' : diameter of the directed graph				
D'	59075	102151	N/A	N/A	D : diameter of the undirected or				
D	50500	54843	5527	7849	symmetrized graph				

Table 1: Tested graphs. Since it is hard to obtain the exact value of D and D' , the number shown is a lower bound obtained by at least 1000 sampled searches on each graph.

and is a tunable parameter. In this way, VGC 1) greatly reduces the number of rounds, since each round may advance multiple hops from the current frontier, and 2) quickly accumulates a large frontier size since the next frontier contains multiple-hop neighbors from the current frontier, and thus yields sufficient parallel tasks throughout the algorithm. Both outcomes effectively reduce (or hide) synchronization costs and allow for much better parallelism.

Data Structure Design. Another useful technique for the SCC algorithm is a data structure called *hash bag* [24]. It maintains a dynamic set of vertices as the frontiers for parallel graph algorithms. For page limit, we refer the readers to [24] for more details.

2.2 Other Algorithms

Parallel SSSP. The SSSP algorithm in PASGAL is based on the *stepping algorithm framework* [11], and uses VGC and hash bags introduced in Sec. 2.1 to accelerate the frontier traversing.

Parallel BFS. Using VGC, we implemented a new BFS algorithm in PASGAL. We also use hash bags to maintain the frontiers. Our BFS algorithm is similar to SSSP where the output distance is the hop distance from the source. For any vertex encountered in a local search from vertex v , we add it to the corresponding frontier if its hop distance can be updated by v . Note that due to local search, a vertex can be visited multiple times instead of exactly once as in standard BFS. This is because the distance of a vertex v updated by a local search is not necessarily the shortest distance of v , leading to extra work. To reduce this overhead, one special technique for BFS is that we maintain multiple frontiers, where frontier i maintains vertices with distance 2^i from the current frontier. In this way, we obtain the benefit of BFS by exploring multiple hops in one round, and also avoid visiting too many “unready” vertices in the frontier. We also use the direction optimization [4] to improve performance.

Parallel Biconnectivity. Different from other problems, the major performance gain of the BCC algorithm in PASGAL is due to algorithm redesign to achieve *stronger theoretical bounds*. The performance bottleneck of previous BCC algorithms either comes from the use of BFS that requires $O(D)$ rounds of global synchroniza-




Figure 2: Speedup of parallel algorithms over the standard sequential algorithm. y -axis is in log-scale. Bars below 1.0 mean the parallel algorithm is slower than a sequential one. Some bars are invisible because they are close to 1. “N/A”: not applicable. “OOM”: out-of-memory.

tions (e.g., GBBS [9]), or requires $O(m)$ auxiliary space and does not scale to large graphs (e.g., Tarjan-Vishkin [22]). PASGAL uses the FAST-BCC algorithm in [12]. By redesigning the algorithm, FAST-BCC avoids the use of BFS, and achieves $O(n + m)$ work, polylogarithmic span, and $O(n)$ auxiliary space. We also use VGC and hash bags to further improve the performance.

3 Experimental Results

Library Design. We release the code of PASGAL [10]. PASGAL is implemented in C++ using ParlayLib [6] for fork-join parallelism and some parallel primitives (e.g., sorting). The four algorithms (BFS, BCC, SCC and SSSP) are provided in the subdirectories. A readme file about compiling and running the library is provided in the repository. PASGAL supports two different graph formats: the binary format (.bin) from GBBS [9], and the adjacency graph format (.adj) from the Problem-Based Benchmark Suite (PBBS) [2]. **Setup.** We run our experiments on a 96-core (192 hyperthreads) machine with four Intel Xeon Gold 6252 CPUs and 1.5 TB of main memory. We use `numactl -i all` in parallel experiments.

We tested on 22 graphs, including social networks, web graphs, road graphs, k -NN graphs, and synthetic graphs. All the graphs are from existing research papers and public datasets. The graph information is given in Alg. 1. For page limit, we provide the full graph information and corresponding citations in the Appendix. We call the social and web graphs *low-diameter graphs* as they have diameters mostly within a few hundred. We call the road, k -NN, and synthetic graphs *large-diameter graphs* as their diameters are mostly more than a thousand. We symmetrize directed graphs for testing BCC. SCC does not apply to undirected graphs.

We present the performance comparison in Fig. 2. For page limit, we only show results for SCC, BCC, and BFS. For each problem, we also implement the standard sequential algorithm as the baseline, which is a queue-based solution for BFS, Tarjan’s algorithm for SCC [21], and the Hopcroft-Tarjan algorithm for BCC [14]. We compute the relative speedup for each parallel implementation over the sequential algorithm, and present them in Fig. 2. The y -axis is in log-scale. Bars below 1 mean that the parallel algorithm is *slower than the sequential implementation*. The baselines include state-of-the-art graph libraries and implementations, including GBBS [9], GAPBS [5], Tarjan-Vishkin from [12], and Multistep [20].

For all tested problems and all graphs, PASGAL is always compet-

itive on small-diameter graphs: across all graphs, PASGAL is within $1.3\times$ of the running time compared to the fastest baseline on BCC, $2\times$ on BFS, and always faster than all baselines on SCC. Parallel BFS on social networks is one of the most well-studied parallel graph algorithms, and all parallel algorithms achieve superlinear speedup on some social networks due to various optimizations (e.g., the direction optimization [4]). PASGAL achieves good scalability and is $49\text{--}570\times$ faster than the standard sequential algorithm using 192 threads. An interesting future direction is to further make the BFS performance of PASGAL match the best parallel implementation.

On large-diameter graphs, PASGAL achieves much better performance than *all baselines*. On BCC, due to theoretical efficiency, PASGAL consistently outperforms the sequential Hopcroft-Tarjan algorithm on all graphs. It is up to $3.45\times$ faster than the best baseline on each graph. On SCC and BFS, PASGAL is always faster than the sequential baseline except for one graph CH5, which has very large diameter compared to its small size. Different from BCC, our SCC and BFS algorithms do not have strong span bound. Using VGC can only alleviate the scalability issue on large-diameter graphs, but may still be unable to eliminate the issue on adversarial graphs (e.g., a chain). Still, on most real-world large-diameter graphs, PASGAL is almost always the fastest among all parallel implementations, and is up to $5\times$ faster than the best baseline on BFS, and up to $46\times$ on SCC.

4 Conclusion and Future Work

In this paper, we present PASGAL, a scalable graph library specially designed to improve performance on large-diameter graphs. As mentioned, some interesting future directions include further seeking new ideas to improve the performance of BFS on small-diameter graphs that also work well with VGC, as well as further improving the performance for BFS and SCC on very large-diameter graphs. In addition, we believe the techniques in current PASGAL can be extended to more problems, including k -core and other peeling algorithms, k -connectivity, point-to-point shortest paths, etc. We plan to add them to PASGAL in the future.

Acknowledgement

This work is supported by NSF grants CCF-2103483, CCF-2238358, CCF-2339310, and IIS-2227669, the UCR Regents Faculty Development Award, and the Google Research Scholar Program.

References

- [1] 2010. OpenStreetMap © OpenStreetMap contributors. <https://www.openstreetmap.org/>.
- [2] Daniel Anderson, Guy E Blelloch, Laxman Dhulipala, Magdalen Dobson, and Yihan Sun. 2022. The problem-based benchmark suite (PBBS), V2. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*. 445–447.
- [3] Lars Backstrom, Dan Huttenlocher, Jon Kleinberg, and Xiangyang Lan. 2006. Group formation in large social networks: membership, growth, and evolution. In *ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*. 44–54.
- [4] Scott Beamer, Krste Asanović, and David Patterson. 2012. Direction-optimizing breadth-first search. In *International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*. 1–10.
- [5] Scott Beamer, Krste Asanović, and David Patterson. 2015. The GAP benchmark suite. *arXiv preprint arXiv:1508.03619* (2015).
- [6] Guy E. Blelloch, Daniel Anderson, and Laxman Dhulipala. 2020. ParlayLib — a toolkit for parallel algorithms on shared-memory multicore machines. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 507–509.
- [7] Paolo Boldi, Marco Rosa, Massimo Santini, and Sebastiano Vigna. 2011. Layered Label Propagation: A MultiResolution Coordinate-Free Ordering for Compressing Social Networks. In *www*. 587–596.
- [8] Paolo Boldi and Sebastiano Vigna. 2004. The WebGraph Framework I: Compression Techniques. In *www*. ACM Press, Manhattan, USA, 595–601.
- [9] Laxman Dhulipala, Guy E. Blelloch, and Julian Shun. 2021. Theoretically efficient parallel graph algorithms can be fast and scalable. *ACM Transactions on Parallel Computing (TOPC)* 8, 1 (2021), 1–70.
- [10] Xiaojun Dong, Yan Gu, Yihan Sun, and Letong Wang. 2024. PASGAL: Parallel And Scalable Graph Algorithm Library. <https://github.com/ucparlay/PASGAL>.
- [11] Xiaojun Dong, Yan Gu, Yihan Sun, and Yunming Zhang. 2021. Efficient Stepping Algorithms and Implementations for Parallel Shortest Paths. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 184–197.
- [12] Xiaojun Dong, Letong Wang, Yan Gu, and Yihan Sun. 2023. Provably Fast and Space-Efficient Parallel Biconnectivity. In *ACM Symposium on Principles and Practice of Parallel Programming (PPOPP)*. 52–65.
- [13] Jordi Fonollosa, Sadique Sheik, Ramón Huerta, and Santiago Marco. 2015. Reservoir computing compensates slow response of chemosensor arrays exposed to fast varying gas concentrations in continuous monitoring. *Sensors and Actuators B: Chemical* 215 (2015), 618–629.
- [14] John Hopcroft and Robert Tarjan. 1973. Algorithm 447: efficient algorithms for graph manipulation. *Commun. ACM* 16, 6 (1973), 372–378.
- [15] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a social network or a news media?. In *International World Wide Web Conference (WWW)*. 591–600.
- [16] YongChul Kwon, Dylan Nunley, Jeffrey P Gardner, Magdalena Balazinska, Bill Howe, and Sarah Loebman. 2010. Scalable clustering algorithm for N-body simulations in a shared-nothing cluster. In *International Conference on Scientific and Statistical Database Management*. Springer, 132–150.
- [17] Robert Meusel, Oliver Lehmberg, Christian Bizer, and Sebastiano Vigna. 2014. Web Data Commons — Hyperlink Graphs. <http://webdatacommons.org/hyperlinkgraph>.
- [18] Veronica Red, Eric D Kelsic, Peter J Mucha, and Mason A Porter. 2011. Comparing community structure to characteristics in online collegiate social networks. *SIAM review* 53, 3 (2011), 526–543.
- [19] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *aaai*. <https://networkrepository.com>
- [20] George M. Slota, Sivasankaran Rajamanickam, and Kamesh Madduri. 2014. BFS and coloring-based parallel algorithms for strongly connected components and related problems. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 550–559.
- [21] Robert Tarjan. 1972. Depth-first search and linear graph algorithms. *SIAM J. on Computing* 1, 2 (1972), 146–160.
- [22] Robert E Tarjan and Uzi Vishkin. 1985. An efficient parallel biconnectivity algorithm. *SIAM J. on Computing* 14, 4 (1985), 862–874.
- [23] Amanda L Traud, Peter J Mucha, and Mason A Porter. 2012. Social structure of facebook networks. *sma* 391, 16 (2012), 4165–4180.
- [24] Letong Wang, Xiaojun Dong, Yan Gu, and Yihan Sun. 2023. Parallel Strong Connectivity Based on Faster Reachability. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*.
- [25] Yiqiu Wang, Shangdi Yu, Laxman Dhulipala, Yan Gu, and Julian Shun. 2021. GeoGraph: A Framework for Graph Processing on Geometric Data. *ACM SIGOPS Operating Systems Review* 55, 1 (2021), 38–46.
- [26] Jaewon Yang and Jure Leskovec. 2015. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems* 42, 1 (2015), 181–213.
- [27] Yu Zheng, Like Liu, Longhao Wang, and Xing Xie. 2008. Learning transportation mode from raw gps data for geographic applications on the web. In *International World Wide Web Conference (WWW)*. 247–256.

A More Experimental Results

We present the original running time in our experiments and more detailed graph information here.

Tab. 2 presents the information of all graphs tested in the paper. Tab. 3 to 5 presents the running time of all tested algorithms on biconnected components (BCC), breadth-first search (BFS), and strongly connected components (SCC), respectively.

For each table, the last column is always the standard *sequential* algorithm (noted with “*s*”). In the figure in the main paper, the relative speedup are based on this sequential algorithm. For all running times, lower is better. “o.o.m” means out-of-memory. “n.s.” means no-support.

		n	m'	m	D'	D	Notes
Social	LJ	4,847,571	68,993,773	85,702,474	22	19	soc-LiveJournal1 [3]
	FB	59,216,214	N/A	185,044,024	N/A	22	socfb-konect [18, 19, 23]
	OK	3,072,627	N/A	234,370,166	N/A	9	com-orkut [26]
	TW	41,652,231	1,468,365,182	2,405,026,092	18	22	Twitter [15]
	FS	65,608,366	N/A	3,612,134,270	N/A	37	Friendster [26]
Web	WK	6,625,370	165,206,104	300,331,770	62	9	enwiki-2023 [7, 8]
	SD	89,247,739	2,043,203,933	3,880,015,728	145	35	sd-arc [17]
	CW	978,408,098	42,574,107,469	74,744,358,622	506	254	ClueWeb [17]
	HL14	1,724,573,718	64,422,807,961	124,141,874,032	800	366	Hyperlink14 [17]
	HL12	3,563,602,789	128,736,914,167	225,840,663,232	5279	650	Hyperlink12 [17]
Road	AF	33,493,259	44,773,338	88,929,770	8276	3948	Open Street Map (OSM) Africa [1]
	NA	86,951,513	112,869,708	220,285,922	9337	4835	Open Street Map (OSM) North America [1]
	AS	95,735,401	122,836,037	243,624,688	16660	8794	Open Street Map (OSM) Asia [1]
	EU	130,655,972	168,541,220	332,587,928	11814	4410	Open Street Map (OSM) Europe [1]
k NN	CH5	4,208,261	21,041,305	29,650,038	5683	14479	Chem [13, 25], $k = 5$
	GL5	24,876,978	124,384,890	157,442,032	17268	21601	GeoLife [25, 27], $k = 5$
	GL10	24,876,978	248,769,780	309,743,322	13982	9053	GeoLife [25, 27], $k = 10$
	COS5	321,065,547	1,605,327,735	1,957,750,718	1390	1180	Cosmo50 [16, 25], $k = 5$
Synthetic	REC	100,000,000	297,418,030	400,000,000	59075	50500	$10^3 \times 10^5$ Grid [24]
	SREC	100,000,000	203,785,880	335,787,820	102151	54843	Sampled REC [24]
	TRCE	16,002,413	N/A	47,997,626	N/A	5527	Huge traces [19]
	BBL	21,198,119	N/A	63,580,358	N/A	7849	Huge bubbles [19]

Table 2: Information of all tested graphs. n : #vertices. m : #edges in the undirected or symmetrized graph. m' : # directed edges. D : diameter of the undirected or symmetrized graph. D' : diameter of the directed graph. Undirected graphs are underlined.

	Graph	PASGAL	GBBS	Tarjan-Vishkin	Hopcroft-Tarjan*
Social	LJ	0.112	0.165	0.857	2.38
	FB	0.835	0.825	1.95	8.68
	OK	0.106	0.15	2.46	3.3
	TW	1.44	2.72	24.6	65.6
	FT	3.17	6.26	40.7	177.6
Web	WK	0.175	0.248	3.1	5.58
	SD	3.16	5.47	41.7	127.9
	CW	23.8	39.4	o.o.m.	704.5
	HL14	33.7	49.3	o.o.m.	818.3
	HL12	129.8	102.4	o.o.m.	2216.6
Road	AF	0.859	5.77	1.42	17.1
	NA	2.19	8.19	3.62	44.2
	ASIA	2.46	10.4	4.33	49.1
	EU	3.32	11.4	5.84	69
kNN	CH5	0.141	1.31	0.418	0.465
	GL5	0.501	2.65	1.73	4.97
	GL10	0.558	1.66	3.37	7.36
	COS5	8.28	18.4	26.5	175.8
Synthetic	REC	1.47	28.7	4.81	5.95
	SREC	1.44	26.3	4.32	12.5
	TRCE	0.302	4.52	0.674	3.36
	BBL	0.406	5.62	0.958	5.1
Geometric Mean:					
	Social	0.63	0.913	5.50	9.53
	Web	4.97	7.06	-	112
	Road	1.98	8.64	3.38	20.7
	kNN	0.756	3.21	2.84	4.55
	Synthetic	0.714	11.8	1.91	3.68

Table 3: Running time of all tested algorithms on biconnected components (BCC). The parallel implementations include PASGAL (this paper), GBBS [9], and Tarjan-Vishkin [22] (implementation from [12]). The sequential baseline is the Hopcroft-Tarjan algorithm [14].

	Graph	PASGAL	GBBS	Multistep	Tarjan*
Social	LJ	0.033	0.122	0.117	1.31
	FB		Undirected Graph		
	OK		Undirected Graph		
	TW	0.192	0.332	1.31	30.2
	FT		Undirected Graph		
Web	WK	0.054	0.120	0.185	2.71
	SD	1.06	4.93	1.84	44.5
	CW	12.6	38.8	n.s.	267
	HL14	18.8	63.6	n.s.	326
	HL12	96.8	364.0	n.s.	1620
Road	AF	1.40	50.3	19.9	7.32
	NA	1.64	24.6	12.4	19.2
	ASIA	3.34	129	34.1	20.8
	EU	2.34	76.2	31.8	29.2
kNN	CH5	0.528	8.45	18.0	0.224
	GL5	0.903	11.9	23.7	2.34
	GL10	1.62	17.5	17.7	3.41
	COS5	3.27	14.9	58.0	87.1
Synthetic	REC	0.892	41.3	67.3	7.91
	SREC	2.02	30.2	82.2	6.70
	TRCE		Undirected Graph		
	BBL		Undirected Graph		
Geometric Mean					
	Social	0.080	0.201	0.391	6.30
	Web	4.20	14.0	-	112
	Road	2.06	59.0	22.8	17.1
	kNN	1.26	12.7	25.7	3.53
	Synthetic	1.34	35.3	74.4	7.28

Table 4: Running time of all tested algorithms on strongly connected components (SCC). The parallel implementations include PASGAL (this paper), GBBS [9], and Multistep [20]. Multistep does not support the three largest graph because the number of vertices in CW, HL14, and HL12 are larger than 32-bit integers. The sequential baseline is the Tarjan’s algorithm [21].

	Graph	PASGAL	GBBS	GAPBS	Queue-based*
Social	LJ	0.018	0.010	0.030	1.19
	FB	0.104	0.052	0.104	5.14
	OK	0.012	0.007	0.011	2.19
	TW	0.136	0.077	0.116	30.2
	FT	0.214	0.163	0.173	122
Web	WK	0.026	0.014	0.041	2.67
	SD	0.471	0.366	0.421	43.1
	CW	5.69	4.10	3.44	283.3
	HL14	4.19	3.12	2.83	158
	HL12	22.5	12.2	o.o.m.	759
Road	AF	0.05	0.7	0.2	0.45
	NA	0.31	5.4	1.4	11.0
	ASIA	0.18	5	0.9	3.0
	EU	0.49	8.8	2.3	16.3
kNN	CH5	0.095	0.10	0.3	0.047
	GL5	0.095	0.1	0.1	0.10
	GL10	0.08	0.3	0.5	0.32
	COS5	3.12	1.8	5.0	90.1
Synthetic	REC	1.100	14.7	4.1	5.86
	SREC	1.14	11.2	5.7	4.96
	TRCE	0.134	1.81	0.458	1.84
	BBL	0.185	2.92	0.618	2.77
Geometric Mean					
	Social	0.058	0.034	0.059	8.68
	Web	1.46	0.956	-	112
	Road	0.191	3.67	0.886	3.95
	kNN	0.215	0.287	0.517	0.603
	Synthetic	0.420	5.43	1.60	3.49

Table 5: Running time of all tested algorithms on breadth-first search (BFS). The parallel implementations include PASGAL (this paper), GBBS [9], and GAPBS [5]. The sequential baseline is the standard algorithm based on maintaining the visited vertices in a queue.