

Provably Fast and Space-Efficient Parallel Biconnectivity (PPoPP'23 Best Paper)

RIVERSIDE

Xiaojun DongLetong WangYan GuYihan SunFull Version: https://github.com/ucrparlay/FAST-BCC



Problem Definitions

- Given an undirected graph G = (V, E) with n = |V| vertices and m = |E| edges
- A connected component (CC) is a maximal subset in V such that every two vertices in it are connected by a path



- Two vertices are biconnected if they are connected and remain connected after removing any other single vertex
- A biconnected component (BCC) (or a block) is a maximal subset of biconnected vertices



Ours GBBS SM'14 SEQ Ours GBBS SM'14 SEQ 5.88 4.36 HH5 1.14 1.00 ΥT 3.15 7.01 1.00 n 30.51 19.91 5.66 4.11 0.37 1.00 CH5 OK 1.00 n Social 17.92 11.77 GL2 6.24 1.64 1.00 LJ 1.00 n 2.40 GL5 8.53 1.44 n 1.00 ΤW 34.21 17.42 1.00 NN-> 18.93 10.22 GL10 10.59 1.00 4.31 FT 1.00 39.26 21.23 12.75 4.57 GL15 11.88 1.00 MEAN 1.00 5.91 8.92 5.65 1.00 GL20 11.84 6.88 1.00 GG n n 29.74 16.46 COS5 14.16 6.86 SD 1.00 1.00 n n MEAN 8.68 2.42 CW 1.00 1.00 Web 30.37 17.52 n 1.59 10.56 HL14 32.46 19.96 1.00 SQR 18.50 1.00 33.99 29.15 REC 12.48 3.02 HL12 1.00 0.36 1.00 n Synthetic MEAN 24.53 15.68 1.00 SQR' 8.06 0.85 1.00 n n 1.00 REC' 0.48 5.15 0.55 7.81 CA 1.00 Road 6.69 0.60 Chn7 11.97 0.49 1.00 0.04 USA 0.08 1.00 GE 10.77 1.00 11.97 1.43 2.44 Chn8 0.04 0.06 1.00 MEAN MEAN 0.73 11.30 7.18 1.21 1.00 0.27 1.00 0.18 TOTAL MEAN 12.89 2.50 0.96 1.00 n = no support

5 1 2 4 8 16 32 >32 MEAN = geometric mean

The heatmap shows relative speedups for parallel BCC algorithms over the sequential Hopcroft-Tarjan algorithm (SEQ) using 96 cores (192 hyper-threads)
On large-diameter graphs, GBBS and SM'14 only achieve 0.18-2.42x speedup. Can be slower than SEQ!
FAST-BCC achieves the best performance on all graphs
On average, it is significantly faster than all baselines!
About 13x faster than SEQ
About 5x faster than GBBS (previous best implementation)

Experimental Results



"Skeleton" G' = (V', E')

Connectivity of G'

(some postprocessing may be needed)

- Many BCC algorithms follow the skeleton-connectivity framework
- Challenge: generating the skeleton (time-)efficiently and storing the skeleton (space-)efficiently!

Related Work

Algorithms	Generating Skeleton	Storing Skeleton
Sequential Hopcroft-Tarjan [HT73]	<pre>DFS Tree</pre>	Not stored (processed on- the-fly) ✓ Space-efficient
Parallel Tarjan-Vishkin [TV85]	Arbitrary Spanning Tree ✓ Work- and span-efficient!	Skeleton needs $O(m)$ extra space × Space-inefficient!
Existing Parallel Implementations [CB05, DBS18, SM14]	 BFS Tree (Span proportional to graph diameters!) * Low parallelism in the worst case 	 O(n) space, or maintained implicitly in O(n) space ✓ Space-efficient



scalability on all graphs

References

Guojing Cong and David Bader. 2005. An experimental study of parallel biconnected components algorithms on symmetric multiprocessors (SMPs). In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE.
 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.

[3] Xiaojun Dong, Letong Wang, Yan Gu, and Yihan Sun. 2023. Provably Fast and Space-Efficient Parallel Biconnectivity. ACM Symposium on Principles and Practice of Parallel Programming (PPOPP) (2023), 52–65.

[4] John Hopcroft and Robert Tarjan. 1973. Algorithm 447: efficient algorithms for graph manipulation. *Commun. ACM* 16, 6 (1973), 372–378.
[5] George Slota and Kamesh Madduri. 2014. Simple parallel biconnectivity algorithms for multicore platforms. In *IEEE International Conference on High Performance Computing (HiPC)*. IEEE, 1–10.
[6] Robert E Tarjan and Uzi Vishkin. 1985. An efficient parallel biconnectivity algorithm. *SIAM J. on Computing* 14, 4 (1985), 862–874.

Our algorithm (FAST-BCC) Arbitrary Spanning Tree

✓ Work- and spanefficient! Maintained implicitly using
O(n) space
✓ Space-efficient

Algorithm Overview			
Input Graph G: contains 3 BCCs {s, u}, {r, s, t, v, w, x}, {t, y, z}.	Step 2: Rooting. Generate rooted spanning trees. Step 2.1: Based on the spanning tree T of G. Create the linked list of the Euler tour of T. U V Step 2: Rooting. Generate rooted spanning trees. Step 2.1: Based on the spanning tree T of G. Create the linked list of the Euler tour of T. U V Step 2.2: Set any vertex as the root and run list ranking. The result implies the tree edge directions. V Step 2.2: Set any vertex as the root and run list ranking. The result implies the tree edge directions. V Step 2.2: Set any vertex as the root and run list ranking. The result implies the tree edge directions. V Step 2.2: Set any vertex as the root and run list ranking. The result implies the tree edge directions.	Step 3: Tagging. Compute tags (first/last/low/high/) for each vertex. Use these tags to identify fence, plain, cross, and back edges. Tree edge → Fence edge → Plain edge Non-tree edge — Back edge — Back edge — Cross edge	
Step 4: Last CC. Run CC on the ske Step 4.1: Find the CCs of the skeleton <i>G</i> ' only with cross and plain edges in <i>G</i> (solid edges in Step 3). Ignore the root.	eton. () () () () () () () () () ()	S t Y V V Y Z BCC1 {s, u}: Head s + {u} BCC2 {s, t, v, w, x}: Head r + {s, t, v, w, x}: Head r + {s, t, v, w, x} BCC3 {t, y, z}: Head t + {y, z}	