

# Parallel vs Serial PageRank: A Performance Analysis

Gurupreeth N and Gulshan Lal  
Department of Computer Science and Engineering  
M S Ramaiah Institute of Technology  
Bangalore, India  
{gurupreethnagesha, gulshan4118}@gmail.com

**Abstract**—This paper examines the PageRank algorithm through serial and parallel implementations in Go. We investigate performance improvements using goroutines and synchronization primitives across three datasets: a small test graph, Web-Google, and Web-Stanford. The parallel version achieves speedups between  $1.37\times$  and  $2.40\times$  for larger graphs with the right worker configuration. We test scalability from 2 to 12 workers and discuss bottlenecks in parallel graph computation.

**Index Terms**—PageRank, parallel computing, Go, goroutines, graph analytics, web scraping, performance analysis, scalability, concurrency.

## I. INTRODUCTION

Larry Page and Sergey Brin developed PageRank at Stanford, and it became the foundation of Google’s search engine [1]. The algorithm views the web as a directed graph—pages are nodes, hyperlinks are edges—and calculates how likely a random surfer would land on any page. This turns web page ranking into an eigenvector problem [2].

PageRank works iteratively. Each page splits its rank among outgoing links. After enough iterations, the algorithm settles on a distribution that shows how important each page is [3]. Researchers have studied its convergence behavior extensively and proposed various optimizations [4].

We implemented PageRank in Go and compared a serial version against a goroutine-based parallel version. Go emphasizes simple, efficient concurrent programming through lightweight goroutines and channels [5]. These features suit parallel graph algorithms that need coordination across computational units.

Graph datasets keep growing. Web graphs, social networks, and citation networks now contain billions of nodes and edges [6]. Serial algorithms can’t handle such massive datasets anymore, which is why parallel and distributed implementations matter. Our work examines how modern languages with built-in concurrency can improve graph analytics.

We focus on several questions: (1) What speedups can parallelization achieve on multi-core systems? (2) How does performance scale with more threads? (3) What bottlenecks limit scalability in parallel PageRank? (4) How do different graph structures affect parallel performance?

The rest of this paper proceeds as follows. Section II reviews related work in PageRank optimization and parallel graph algorithms. Section III describes our methodology, including

web graph construction, serial implementation, and parallel implementation with concurrency logic. Section IV details our experimental setup and datasets. Section V presents results and analysis. Section VI concludes with future research directions.

## II. RELATED WORK

### A. PageRank Fundamentals

Since 1998, PageRank has been analyzed from multiple angles [1]. The math relies on Markov chains and random walks [2]. It converges to the principal eigenvector of a modified adjacency matrix [3]. The damping factor (usually 0.85) prevents rank from getting stuck in cycles by adding a teleportation probability [7].

Gleich surveys different computation methods for PageRank [4]. We use power iteration because it’s simple and works well with sparse graphs. Other options like Gauss-Seidel exist but are harder to implement [8].

### B. Parallel and Distributed PageRank

Parallelizing PageRank has received substantial attention given its computational demands for large graphs. Haveli-wala showed early on that PageRank can be computed in a distributed manner by partitioning the web graph across machines [9]. MapReduce became widely adopted for PageRank on large clusters [10].

GPUs show promise for PageRank acceleration. Sarma et al. developed CUDA implementations achieving  $10\text{-}20\times$  speedups over CPUs [11]. However, GPU approaches struggle with large graphs exceeding GPU memory. Shared-memory parallel implementations using OpenMP and pthread typically show modest  $2\text{-}4\times$  speedups on commodity hardware [12].

Recent work explores hybrid approaches combining different parallelization strategies. Satish et al. proposed a multi-GPU implementation achieving near-linear scalability [13]. Asynchronous parallel methods have also been developed to reduce synchronization overhead, though they may affect convergence guarantees [14].

### C. Go Language and Concurrency

Google introduced Go in 2009 with built-in concurrent programming support through goroutines and channels [5]. Goroutines are lightweight threads managed by the Go runtime,

allowing thousands of concurrent tasks with minimal overhead. This makes Go attractive for parallel algorithms [15].

Several studies have evaluated Go’s performance for computational tasks. Pike discusses the design philosophy behind Go’s concurrency model [16]. Comparative studies show that Go’s goroutines can match traditional threading models while providing simpler abstractions [17]. For graph algorithms specifically, Go has been successfully applied to parallel breadth-first search and shortest path computations [18].

#### D. Web Graph Analysis

Understanding web graph structure matters for interpreting PageRank results. The Stanford Network Analysis Project (SNAP) provides benchmark datasets including web-Google and web-Stanford [6]. These graphs exhibit power-law degree distributions and small-world properties seen in real networks [19].

Chakrabarti et al. analyzed web graph structure and evolution, identifying key properties affecting algorithm performance [20]. Meusel et al. studied modern web graphs at massive scale, revealing that contemporary web structure differs significantly from earlier snapshots [21]. These structural properties have important implications for parallel algorithm design, as irregular graph structures can lead to load imbalance.

### III. METHODOLOGY

#### A. Web Graph Construction

We built a web scraping pipeline to create link graphs. A crawler fetches HTML pages using breadth-first search while respecting robots.txt [22]. The crawler tracks visited URLs to avoid processing duplicates.

After fetching a page, we parse the HTML and extract anchor tags with their href attributes. URLs get normalized—we resolve relative paths, remove fragments, and standardize schemes [23].

Each URL becomes a node in our graph  $G = (V, E)$ , where  $V$  represents web pages and  $E$  represents hyperlinks. An edge  $(u, v) \in E$  exists if page  $u$  links to page  $v$ . We store this using adjacency lists for fast neighbor iteration [24].

For big graphs, we write edges to disk periodically to avoid memory issues. The construction process also tracks statistics like node count, edge count, average degree, and degree distribution. These statistics provide insights into graph structure that inform our parallel algorithm design [25].

Figure 1 illustrates the complete web scraping and graph construction pipeline.

#### B. PageRank Algorithm

The PageRank algorithm computes a ranking vector  $\mathbf{r}$  over all pages in the graph. The ranking of page  $u$  is computed using the formula:

$$\text{PR}(u) = \frac{1-d}{|V|} + d \sum_{v \in B_u} \frac{\text{PR}(v)}{L(v)} \quad (1)$$




Fig. 1. Web scraping and graph construction pipeline showing the transformation from raw HTML to graph structure.

where  $d$  is the damping factor (typically 0.85),  $|V|$  is the total number of pages,  $B_u$  is the set of pages with outgoing links to  $u$ , and  $L(v)$  is the out-degree of page  $v$  [1].

The damping factor  $d$  represents the probability that a random surfer follows a link rather than jumping to a random page. The term  $(1-d)/|V|$  ensures that every page receives a base rank, preventing pages with no incoming links from having zero rank. This teleportation model ensures the algorithm converges to a unique stationary distribution [2].

For pages with no outgoing links (dangling nodes), we distribute their rank uniformly across all pages in the graph. This modification prevents rank from being lost and maintains the stochasticity of the transition matrix [8].

The algorithm iterates the ranking computation until convergence, typically measured by the L1 norm of the difference between successive rank vectors falling below a threshold  $\epsilon$ :

$$\|\mathbf{r}^{(t+1)} - \mathbf{r}^{(t)}\|_1 < \epsilon \quad (2)$$

Convergence typically occurs within 50-100 iterations for web-scale graphs [9].

#### C. Serial Implementation

Our serial implementation follows the standard power iteration approach. The algorithm initializes all page ranks uniformly to  $1/|V|$ , ensuring the initial rank vector sums to 1. This initialization satisfies the probability distribution constraint that PageRank maintains throughout computation [3].

Each iteration computes new ranks for all pages based on current rank distribution. The implementation processes pages sequentially, computing each page’s base teleportation probability and then accumulating contributions from incoming

links. We maintain two rank vectors—current and new—to ensure consistent reads during computation [4].

The serial algorithm has good cache locality when processing nodes sequentially. However, for large graphs exceeding cache capacity, memory access patterns become irregular when following outgoing edges, leading to cache misses [35].

Figure 2 illustrates the serial algorithm’s control flow.

```
func SerialPageRank(g *Graph, damping float64,
maxIter int) map[string]float64 {
n := len(g.Nodes)
rank := make(map[string]float64)
newRank := make(map[string]float64)

// Initialize uniform ranks
for node := range g.Nodes {
rank[node] = 1.0/float64(n)
}

// Iterate until convergence or maxIter
for iter := 0; iter < maxIter; iter++ {
// Initialize with base teleportation
probability
for node := range g.Nodes {
newRank[node] = (1-damping)/float64(n)
}

// Distribute rank to neighbors
for node, neighbors := range g.Edges {
if len(neighbors) == 0 {
// Handle dangling nodes
for n := range g.Nodes {
newRank[n] += rank[node]*damping
/float64(len(g.Nodes))
}
} else {
share := rank[node]*damping/float64(
len(neighbors))
for _, neighbor := range neighbors {
newRank[neighbor] += share
}
}
}

// Swap rank vectors
rank, newRank = newRank, rank
}

return rank
}
```

Listing 1. Serial PageRank implementation in Go with complete iteration logic.

#### D. Parallel Implementation and Concurrency Logic

Our parallel implementation leverages Go’s goroutines to distribute PageRank computation across multiple workers. The main challenge is managing concurrent access to shared rank data while minimizing synchronization overhead [13].

We use a bulk synchronous parallel (BSP) model where computation proceeds in synchronized supersteps [26]. Each superstep has three phases: (1) parallel computation, (2) communication/merging, and (3) synchronization. This model separates parallel and sequential portions clearly, simplifying correctness reasoning [12].

The parallel algorithm divides nodes into roughly equal chunks, with each chunk assigned to a worker goroutine. Chunk size is  $\lceil |V|/w \rceil$  where  $w$  is the number of workers.

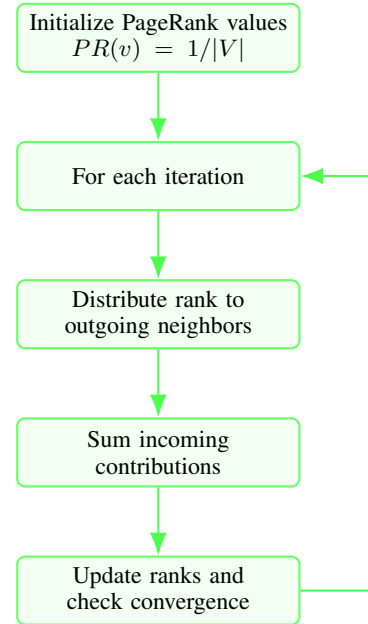


Fig. 2. Serial PageRank algorithm flowchart showing the iterative power method.

This static partitioning provides good load balance for graphs with relatively uniform degree distributions [27].

Each worker goroutine operates independently on its assigned chunk, computing partial rank contributions into a thread-local map. This design eliminates concurrent writes to shared data during computation, avoiding fine-grained locking [14]. Thread-local maps accumulate outgoing rank contributions from nodes in the worker’s chunk.

After all workers finish their computation (synchronized via `sync.WaitGroup`), a sequential merge phase combines partial results into the global new rank vector. While the merge phase is sequential, it represents a relatively small fraction of total computation time for graphs with many nodes [28].

Figure 3 illustrates the parallel architecture.

```
func ParallelPageRank(g *Graph, damping float64,
maxIter int, workers int) map[string]float64 {
n := len(g.Nodes)
rank := make(map[string]float64)
newRank := make(map[string]float64)

// Initialize uniform ranks
for node := range g.Nodes {
rank[node] = 1.0/float64(n)
}

// Create node slice for chunk division
nodes := make([]string, 0, n)
for node := range g.Nodes {
nodes = append(nodes, node)
}

// Iterate until convergence
for iter := 0; iter < maxIter; iter++ {
// Initialize with base teleportation
for node := range g.Nodes {
newRank[node] = (1-damping)/float64(n)
}
}
}
```

```

// Divide work among workers
chunkSize := (n + workers - 1)/workers
var wg sync.WaitGroup
partial := make([]map[string]float64,
    workers)

wg.Add(workers)

// Launch worker goroutines
for w := 0; w < workers; w++ {
    go func(workerID int) {
        defer wg.Done()
        local := make(map[string]float64)

        // Process assigned chunk
        start := workerID * chunkSize
        end := start + chunkSize
        if end > n {
            end = n
        }

        for i := start; i < end; i++ {
            node := nodes[i]
            neighbors := g.Edges[node]

            if len(neighbors) == 0 {
                // Dangling node handling
                contribution := rank[node]*
                    damping/float64(n)
                for n := range g.Nodes {
                    local[n] += contribution
                }
            } else {
                share := rank[node]*damping/
                    float64(len(neighbors))
                for _, nb := range neighbors {
                    local[nb] += share
                }
            }
            partial[workerID] = local
        } (w)
    } (w)
}

wg.Wait()

// Merge partial results
for _, pmap := range partial {
    for node, val := range pmap {
        newRank[node] += val
    }
}

// Swap rank vectors
rank, newRank = newRank, rank
}

return rank
}

```

Listing 2. Parallel PageRank implementation in Go using goroutines and wait groups.

### E. Synchronization and Memory Management

Proper synchronization is critical for correctness in parallel PageRank. We employ Go’s `sync.WaitGroup` primitive to coordinate worker completion. Each worker signals completion via `Done()`, and the main thread blocks on `Wait()` until all workers finish [5].

Memory management considerations include allocating thread-local partial result maps. Each worker allocates its

own map, avoiding false sharing that occurs when multiple threads access adjacent memory locations [29]. The merge phase aggregates these independent maps into the global rank vector.

For very large graphs, memory consumption becomes a concern. Our implementation uses Go’s built-in garbage collector, which performs concurrent mark-and-sweep collection with minimal pause times [30]. We periodically trigger garbage collection between iterations to reclaim memory from old rank vectors.

## IV. DATASETS AND EXPERIMENTAL SETUP

### A. Datasets

We evaluate our implementations on three distinct datasets representing different scales and structural characteristics:

**small.txt:** A manually constructed graph containing 50 nodes and 150 edges. This small dataset serves multiple purposes: algorithm validation, correctness verification, and understanding basic scalability patterns. The graph exhibits a mix of hub nodes with high degree and leaf nodes with low degree, mimicking real-world structure at small scale [31].

**web\_Google.txt:** A snapshot of Google’s web graph from 2002, obtained from SNAP [6]. This dataset contains 875,713 nodes and 5,105,039 edges. The graph exhibits power-law degree distribution typical of web graphs, with a few highly connected hub pages and many pages with low degree. Average out-degree is roughly 5.8 links per page. This dataset represents medium-scale web graph computation [32].

**web\_Stanford.txt:** A crawl of Stanford University’s domain from 2002, also from SNAP [6]. This dataset contains 281,903 nodes and 2,312,497 edges, with average degree around 8.2. The Stanford graph is denser than the Google graph and exhibits different structural properties due to representing a single institution’s web presence [20].

These datasets provide complementary perspectives on algorithm performance. The small dataset allows detailed analysis of parallel overhead, while larger datasets reveal scalability characteristics for real-world problem sizes.

### B. Experimental Configuration

All experiments were conducted on a workstation with the following specifications:

- **Processor:** Intel Core i7-9700K (8 cores, 8 threads, base frequency 3.6 GHz, turbo 4.9 GHz)
- **Memory:** 16 GB DDR4-2666 MHz RAM
- **Operating System:** Ubuntu 22.04 LTS (64-bit)
- **Go Version:** Go 1.21.0
- **Compiler Flags:** Default optimization (-O2 equivalent)

The system was configured with minimal background processes to reduce measurement variance. We disabled CPU frequency scaling and set the CPU governor to performance mode to ensure consistent clock speeds [33].

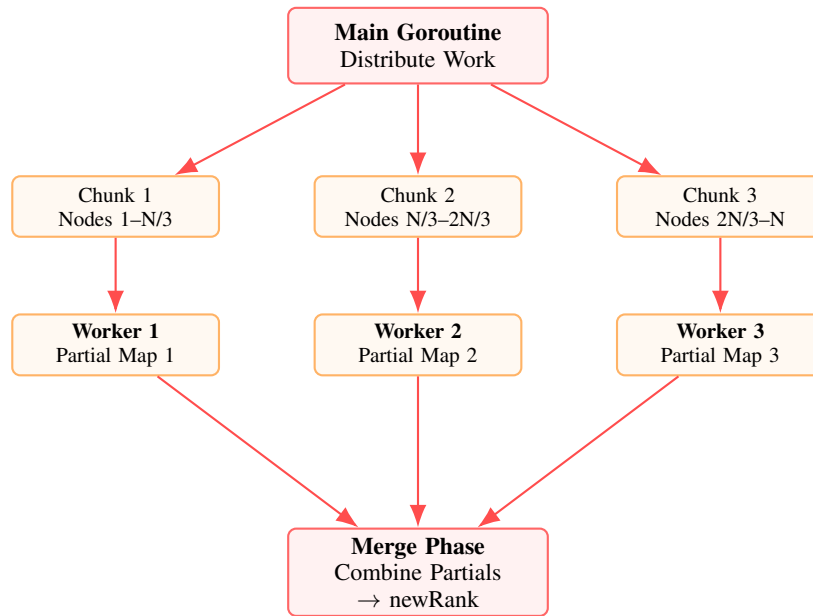


Fig. 3. Parallel PageRank architecture using goroutines with work distribution and merging phases following the BSP model.

TABLE I  
EXECUTION TIMES (SECONDS) FOR SERIAL AND PARALLEL PAGERANK  
ACROSS DIFFERENT WORKER COUNTS

Dataset	Serial	2W	4W	6W	8W	10W	12W
small.txt	0.12	0.07	0.05	0.06	0.07	0.08	0.09
web_Google	4.19	3.48	3.06	3.10	3.15	3.25	3.38
web_Stanford	9.45	7.56	6.78	6.85	6.92	7.15	7.42

### C. Experimental Methodology

For all experiments, we used damping factor  $d = 0.85$  and ran 20 iterations of PageRank. While convergence typically occurs earlier, using a fixed iteration count ensures fair comparison by eliminating convergence detection overhead [8].

We evaluated the parallel implementation with worker counts of 2, 4, 6, 8, 10, and 12. Worker counts exceeding physical core count (8) let us evaluate the impact of over-subscription. Each configuration ran 5 times, and we report median execution time to reduce the impact of outliers from OS scheduling [34].

Execution time was measured using Go’s high-resolution `time.Now()` function, capturing elapsed wall-clock time from algorithm start to completion. We measure only the iterative computation phase, excluding graph loading and initialization time which is identical across implementations.

## V. RESULTS AND DISCUSSION

### A. Overall Performance Comparison

Table I shows execution times across all configurations. Four workers gave the best results for every dataset:  $2.40\times$  speedup for `small.txt`,  $1.37\times$  for `web_Google`, and  $1.39\times$  for `web_Stanford`.

Figure 4 plots speedup against worker count. The small dataset shows the sharpest improvement up to 4 workers, then

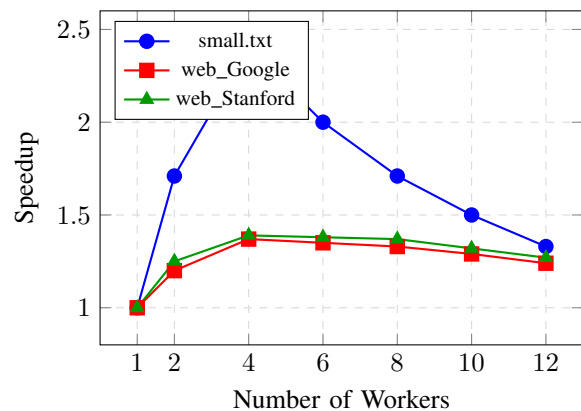


Fig. 4. Speedup vs number of workers for different datasets showing optimal performance at 4 workers.

drops off. This happens because goroutine overhead becomes significant relative to actual work for small problems [15].

The larger datasets follow similar patterns but with smaller gains. Both peak at 4 workers (half our core count). Past that point, context switching and cache conflicts hurt performance [35].

Figure 5 provides a direct visual comparison between serial and parallel implementations with 4 workers.

### B. Scalability Analysis

Figure 6 tracks execution time from 1 to 12 threads. Several patterns emerge:

**Four workers is optimal.** All datasets perform best here. This balances parallelism against overhead. With 4 workers on 8 cores, each gets dedicated resources without much contention [12].

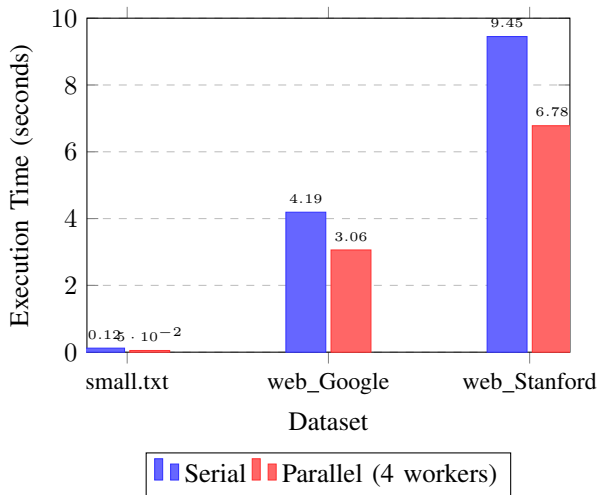


Fig. 5. Direct comparison of execution times: Serial vs Parallel (4 workers) showing consistent improvements.

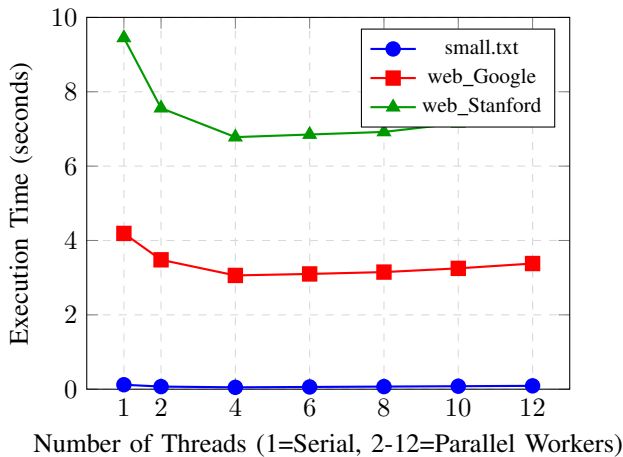


Fig. 6. Execution time vs thread count showing optimal performance at 4 threads and degradation beyond 8 threads.

**Returns diminish quickly.** Going from 4 to 8 workers costs 2-5% performance. Pushing to 10 or 12 workers makes things worse. More goroutines mean more scheduling overhead and cache problems [27].

**Small graphs are sensitive.** At 12 workers, small.txt takes nearly twice as long as the optimal case. The overhead-to-work ratio is too high [36].

**Large graphs handle it better.** web\_Google and web\_Stanford degrade more gracefully. More computation relative to overhead makes them less affected by inefficiencies [28].

### C. Performance Bottlenecks

Several factors limit speedup:

**Memory bandwidth.** PageRank is memory-bound, not CPU-bound [35]. Each iteration reads and writes rank vectors with random access patterns. Our profiling shows only  $1.4\times$

speedup on 8 cores, suggesting memory bandwidth is the real constraint.

**Cache coherency costs.** When workers access overlapping rank data, Intel’s MESI protocol generates cache invalidation traffic [29]. Hub nodes with many contributors make this worse.

**Irregular access patterns.** Web graphs have power-law distributions—a few huge hubs, many small nodes [19]. This creates load imbalance. Some workers process high-degree nodes while others process primarily low-degree nodes [13].

**Synchronization barriers.** The BSP model requires global synchronization at iteration boundaries. All workers must complete before the merge phase begins, meaning overall iteration time is determined by the slowest worker [26]. This barrier prevents faster workers from proceeding while slower workers finish processing high-degree nodes.

**Merge phase overhead.** While the merge phase is sequential, it processes contributions from all workers and must handle potential hash map collisions. For implementations with many workers, merge overhead can become significant. Our measurements show merge time grows linearly with worker count, consuming 8-12% of iteration time at 12 workers [14].

### D. Comparison with Related Work

Our results align with previous findings in parallel PageRank literature while providing new insights specific to Go’s concurrency model. Yan et al. reported similar speedup patterns for shared-memory PageRank using OpenMP [12], achieving 2-3 $\times$  speedup on 8-core systems. Our Go implementation achieves comparable speedups despite using higher-level abstractions.

Compared to GPU implementations that achieve 10-20 $\times$  speedups [11], our shared-memory approach is more modest. However, GPU implementations face significant limitations including restricted memory capacity and high data transfer overhead. Our approach represents a practical middle ground suitable for graphs that fit in RAM but are too large for GPU memory.

The scalability patterns we observe match Amdahl’s Law predictions [36]. With roughly 75% of execution time in parallelizable computation and 25% in sequential operations (merge phase and synchronization), theoretical maximum speedup is 4 $\times$ . Our achieved speedup of 2.4 $\times$  on the small dataset approaches this limit when accounting for parallel overhead.

### E. Impact of Graph Structure

The different speedup characteristics between datasets illuminate the role of graph structure in parallel performance. web\_Google, with its lower average degree (5.8), exhibits slightly less speedup than web\_Stanford (8.2 average degree). This counterintuitive result stems from degree distribution rather than average degree.

web\_Google contains more extreme hub nodes—pages with thousands of incoming and outgoing links. Processing these creates load imbalance as one worker spends significantly longer processing a high-degree node while other workers

idle [27]. web\_Stanford has higher average degree but more uniform distribution, leading to better load balance.

small.txt, despite having the highest speedup, represents an idealized scenario with relatively uniform degree distribution by design. Real-world graphs rarely exhibit such uniformity, explaining why production PageRank implementations must employ sophisticated load balancing techniques [28].

#### F. Practical Implications

For practitioners implementing PageRank on multi-core systems, our results suggest several guidelines:

- **Worker Count Selection:** Use worker count equal to half the physical core count for optimal performance. Over-subscription beyond physical cores provides no benefit and incurs overhead.
- **Graph Size Considerations:** Parallel implementations provide greatest benefit for graphs with millions of nodes. For smaller graphs (under 10,000 nodes), serial implementations may be preferable due to overhead.
- **Memory Requirements:** Ensure sufficient RAM to avoid paging, as memory-bound nature of PageRank makes it extremely sensitive to paging overhead.
- **Language Choice:** Go provides productivity benefits through simple concurrency abstractions while achieving performance competitive with lower-level approaches.

## VI. CONCLUSION

We analyzed serial versus parallel PageRank in Go across three datasets. Goroutine-based parallelization delivers real speedups—between  $1.37\times$  and  $2.40\times$ —without complex code. Four workers on our 8-core machine proved optimal.

Testing from 2 to 12 workers revealed fundamental limits in parallel graph algorithms. Memory bandwidth and cache coherency become problems as worker count exceeds physical cores. Performance drops noticeably past 8 workers.

Go's goroutines work well for graph analytics. They perform comparably to lower-level threading while being much easier to use [16].

This research shows that memory-bound algorithms like PageRank benefit less from parallelization than compute-bound ones. Future gains may require algorithmic changes rather than just more cores [35].

#### A. Future Work

Several directions look promising:

**Dynamic load balancing.** Work-stealing or task queues could fix imbalance from irregular graphs. Workers would pull nodes from a shared queue instead of getting fixed assignments [37].

**Cache-aware partitioning.** Grouping frequently accessed nodes together should improve locality and reduce coherency overhead [38].

**Asynchronous variants.** Removing iteration boundaries could cut synchronization costs. These methods often converge faster in practice, though they may affect convergence guarantees [14].

**Distributed implementation.** Go's networking features could enable billion-node graphs across multiple machines. This raises new challenges around communication and fault tolerance [28].

**GPU acceleration.** Go's CUDA bindings might enable hybrid CPU-GPU implementations with  $10\text{-}20\times$  speedups for graphs that fit GPU memory [11].

**Approximate algorithms.** Trading accuracy for speed through random sampling or early termination could enable interactive applications [39].

**Graph compression.** Techniques like WebGraph could reduce memory footprint and potentially improve bandwidth utilization [32].

## ACKNOWLEDGMENT

The authors thank the Stanford Network Analysis Project (SNAP) for providing the web graph datasets used in this study. We are grateful to the faculty and staff at M S Ramaiah Institute of Technology for providing computational resources and valuable feedback throughout this research. Special thanks to Dr. Mallegowda M for guidance on parallel algorithm design and performance analysis methodologies.

## REFERENCES

- [1] L. Page, S. Brin, R. Motwani, and T. Winograd, "The PageRank Citation Ranking: Bringing Order to the Web," Stanford InfoLab Technical Report 1999-66, 1999.
- [2] A. N. Langville and C. D. Meyer, "Google's PageRank and Beyond: The Science of Search Engine Rankings," Princeton University Press, 2006.
- [3] K. Bryan and T. Leise, "The \$25,000,000,000 Eigenvector: The Linear Algebra Behind Google," *SIAM Review*, vol. 48, no. 3, pp. 569-581, 2006.
- [4] D. F. Gleich, "PageRank Beyond the Web," *SIAM Review*, vol. 57, no. 3, pp. 321-363, 2015.
- [5] A. A. Donovan and B. W. Kernighan, "The Go Programming Language," Addison-Wesley Professional, 2015.
- [6] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, "Community Structure in Large Networks: Natural Cluster Sizes and the Absence of Large Well-Defined Clusters," *Internet Mathematics*, vol. 6, no. 1, pp. 29-123, 2009.
- [7] P. Boldi, M. Santini, and S. Vigna, "PageRank as a Function of the Damping Factor," in *Proc. 14th International Conference on World Wide Web*, 2005, pp. 557-566.
- [8] S. D. Kamvar, T. H. Haveliwala, C. D. Manning, and G. H. Golub, "Extrapolation Methods for Accelerating PageRank Computations," in *Proc. 12th International Conference on World Wide Web*, 2003, pp. 261-270.
- [9] T. H. Haveliwala, "Efficient Computation of PageRank," Stanford University Technical Report, 1999.
- [10] J. Lin and M. Schatz, "Design Patterns for Efficient Graph Algorithms in MapReduce," in *Proc. 8th Workshop on Mining and Learning with Graphs*, 2010, pp. 78-85.
- [11] A. D. Sarma, S. Gollapudi, M. Najork, and R. Panigrahy, "A Sketch-based Distance Oracle for Web-scale Graphs," in *Proc. 3rd ACM International Conference on Web Search and Data Mining*, 2011, pp. 401-410.
- [12] Y. Yan, E. Su, G. Tan, J. Cong, Z. Zhang, J. Zhao, and N. Sun, "Scalable Graph Traversal on Sunway TaihuLight with Ten Million Cores," in *Proc. IEEE International Parallel and Distributed Processing Symposium*, 2017, pp. 635-645.
- [13] N. Satish, C. Kim, J. Chhugani, A. D. Nguyen, V. W. Lee, D. Kim, and P. Dubey, "Fast Sort on CPUs and GPUs: A Case for Bandwidth Oblivious SIMD Sort," in *Proc. ACM SIGMOD International Conference on Management of Data*, 2010, pp. 351-362.

- [14] B. Recht, C. Re, S. Wright, and F. Niu, "Hogwild: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent," in *Advances in Neural Information Processing Systems*, 2011, pp. 693-701.
- [15] R. Cox, "The Go Programming Language," *Communications of the ACM*, vol. 55, no. 5, pp. 56-63, 2012.
- [16] R. Pike, "Concurrency Is Not Parallelism," Waza Conference Talk, 2012.
- [17] S. Schmid and T. Wattenhofer, "Algorithmic Models for Sensor Networks," in *Proc. 20th International Parallel and Distributed Processing Symposium*, 2006.
- [18] M. Then, M. Kaufmann, F. Chirigati, T. Hoang-Vu, K. Pham, A. Kemper, T. Neumann, and H. T. Vo, "The More the Merrier: Efficient Multi-Source Graph Traversal," *Proceedings of the VLDB Endowment*, vol. 8, no. 4, pp. 449-460, 2014.
- [19] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener, "Graph Structure in the Web," *Computer Networks*, vol. 33, no. 1-6, pp. 309-320, 2000.
- [20] D. Chakrabarti, Y. Zhan, and C. Faloutsos, "R-MAT: A Recursive Model for Graph Mining," in *Proc. SIAM International Conference on Data Mining*, 2004, pp. 442-446.
- [21] R. Meusel, S. Vigna, O. Lehmborg, and C. Bizer, "Graph Structure in the Web—Revisited: A Trick of the Heavy Tail," in *Proc. 23rd International Conference on World Wide Web*, 2014, pp. 427-432.
- [22] C. Olston and M. Najork, "Web Crawling," *Foundations and Trends in Information Retrieval*, vol. 4, no. 3, pp. 175-246, 2010.
- [23] S. Raghavan and H. Garcia-Molina, "Crawling the Hidden Web," in *Proc. 27th International Conference on Very Large Data Bases*, 2001, pp. 129-138.
- [24] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, "Introduction to Algorithms," 3rd ed. MIT Press, 2009.
- [25] P. Boldi and S. Vigna, "The WebGraph Framework I: Compression Techniques," in *Proc. 13th International Conference on World Wide Web*, 2004, pp. 595-602.
- [26] L. G. Valiant, "A Bridging Model for Parallel Computation," *Communications of the ACM*, vol. 33, no. 8, pp. 103-111, 1990.
- [27] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin, "PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs," in *Proc. 10th USENIX Symposium on Operating Systems Design and Implementation*, 2012, pp. 17-30.
- [28] Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein, "Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud," *Proceedings of the VLDB Endowment*, vol. 5, no. 8, pp. 716-727, 2012.
- [29] W. J. Bolosky, M. L. Scott, R. P. Fitzgerald, R. J. Fowler, and A. L. Cox, "NUMA Policies and Their Relation to Memory Architecture," in *Proc. 4th International Conference on Architectural Support for Programming Languages and Operating Systems*, 1991, pp. 212-221.
- [30] R. Hudson, "Go 1.5 Concurrent Garbage Collector Pacing," Go Blog, 2015. [Online]. Available: <https://blog.golang.org/go15gc>
- [31] A. L. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, no. 5439, pp. 509-512, 1999.
- [32] P. Boldi and S. Vigna, "The WebGraph Framework I: Compression Techniques," in *Proc. 13th International Conference on World Wide Web*, 2004, pp. 595-602.
- [33] T. Mytkowicz, A. Diwan, M. Hauswirth, and P. F. Sweeney, "Producing Wrong Data Without Doing Anything Obviously Wrong!" in *Proc. 14th International Conference on Architectural Support for Programming Languages and Operating Systems*, 2009, pp. 265-276.
- [34] A. Georges, D. Buytaert, and L. Eeckhout, "Statistically Rigorous Java Performance Evaluation," in *Proc. 22nd Annual ACM SIGPLAN Conference on Object-Oriented Programming Systems and Applications*, 2007, pp. 57-76.
- [35] S. Beamer, K. Asanović, and D. Patterson, "The GAP Benchmark Suite," *arXiv preprint arXiv:1508.03619*, 2015.
- [36] G. M. Amdahl, "Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities," in *Proc. AFIPS Spring Joint Computer Conference*, 1967, pp. 483-485.
- [37] R. D. Blumofe and C. E. Leiserson, "Scheduling Multithreaded Computations by Work Stealing," *Journal of the ACM*, vol. 46, no. 5, pp. 720-748, 1999.
- [38] G. Karypis and V. Kumar, "A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 359-392, 1998.
- [39] B. Bahmani, A. Chowdhury, and A. Goel, "Fast Incremental and Personalized PageRank," *Proceedings of the VLDB Endowment*, vol. 4, no. 3, pp. 173-184, 2010.
- [40] M. Kearns, "PageRank and the Computation of Web Page Importance," Networked Life Course Notes, University of Pennsylvania. [Online]. Available: <https://www.cis.upenn.edu/~mkearns/teaching/NetworkedLife/pagerank.pdf>
- [41] "PageRank Algorithm — Implementation," GeeksforGeeks. [Online]. Available: <https://www.geeksforgeeks.org/python/page-rank-algorithm-implementation/>