Scalable Algorithms for Analyzing Large Dynamic Networks using CANDY

Aashish Pandey[1], Arindam Khanda[2], Sriram Srinivasan[3],Sudharshan Srinivasan[3], S M Shovan[2], Farahnaz Hosseini[1], Sajal K. Das[2], Boyana Norris[3],and Sanjukta Bhowmick[1]

[1]University of North Texas, [2]Missouri University of Science and Technology, [3]University of Oregon

Motivation

- Graph queries on large networks leverage the stored graph properties to provide faster results.
- Real-world graphs (social networks, wireless mobile networks) are often dynamic – graph topology and properties change over time.
- Recomputing graph properties from scratch for each change (node or edge insertion/deletion) is computationally inefficient.
- We proposed a generic framework CANDY for developing efficient parallel algorithms to update graph properties on large dynamic networks.
- Using our novel framework, we design parallel algorithms to update Single Source Shortest Path (SSSP), Multi-objective shortest path, Pagerank, and Vertex Color.

CANDY Architecture

- **CANDY**: a parallel, scalable, extendable, and user-friendly software platform for updating important properties of dynamic networks
- Support parallel dynamic network algorithm development on distributed memory, shared memory, and GPUs, and their use through user-friendly interfaces
- Comprehensive cyberinfrastructure supporting innovative research challenges in large-scale, complex, dynamic networks

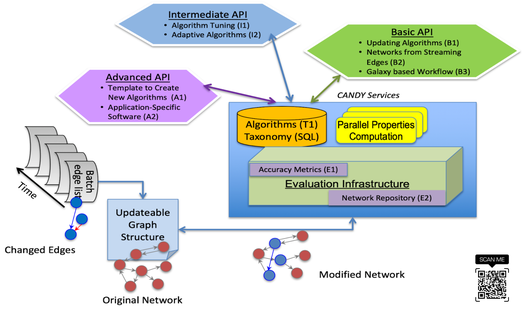

Figure: Overview of CANDY(Cyberinfrastructure for Accelerating Innovation in Network Dynamics)

Property Update Framework

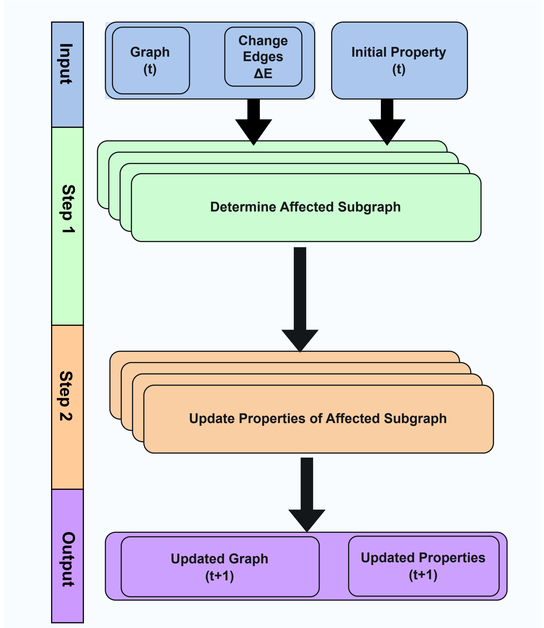

Figure: Parallel Update Framework

Challenges and Intellectual Merits

Challenges

- Efficiently analyzing dynamic networks for real-life applications.
- Designing an efficient algorithm to update with minimal graph traversal.
- Creating parallel graph computation software infrastructure for modern heterogeneous architectures.

Intellectual Merits

- Templates for novel scalable, parallel algorithms for dynamic network analysis
- User-friendly functionality and tools for algorithm creation and modification, accommodating users with various expertise levels.
- Enabling easy access to realistic dynamic graph datasets for efficient evaluation of parallel algorithms.

Parallel Algorithms Developed Using CANDY Framework

Algorithms	*Shared Memory*	*Distributed Memory*	*GPU*
Single Source Shortest Path (SSSP)	✓	✗	✓
Multi-Objective Shortest Path (MOSP)	✓	✗	✗
Strongly-Connected Component(SCC)	✗	✓	✗
Vertex Color Update	✗	✓	✓
Page Rank (PR)	✓	✗	✓

Performance

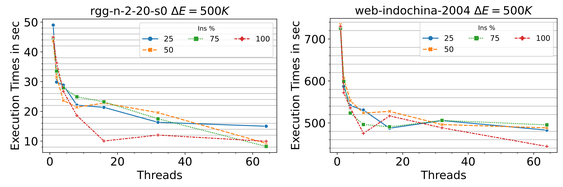

Figure: MOSP update (Shared-mem): Scalability plot

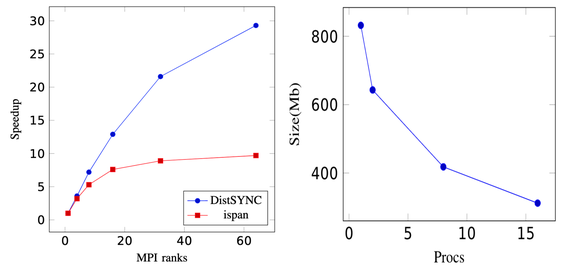

Figure: Scalability(left) and Memory Utilization(right) of SCC update algorithm (RMAT 27)

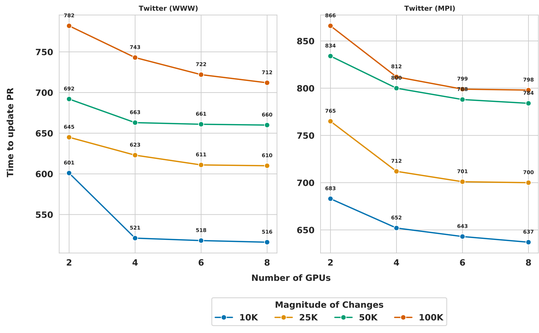

Figure: Multi GPUs PR update

Dynamic Graph Generation

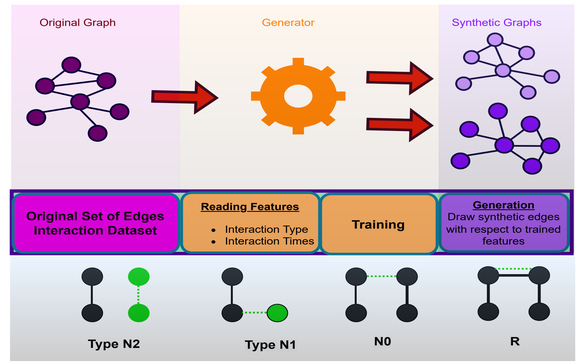

Figure: Dynamic Graph Generator

- Monitoring time and types of interaction during graph evolution.
- Generating synthetic dataset from trained features.

Applications

- Synthesize realistic changes for testing parallel dynamic network algorithms.
- Benchmarks for static snapshot generators.

Conclusion

- Introduced software platform CANDY and parallel update framework.
- Presented parallel algorithms for updating different properties in dynamic networks.

References

- A. Khanda, S. Srinivasan, S. Bhowmick, B. Norris, and S. K. Das, "A parallel algorithm template for updating single-source shortest paths in large-scale dynamic networks," IEEE Transactions on Parallel and Distributed Systems, vol. 33, no. 4, pp. 929–940, 2022.
- S. Srinivasan, S. Pollard, S. K. Das, B. Norris, and S. Bhowmick, "A shared-memory algorithm for updating tree-based properties of large dynamic networks," IEEE Transactions on Big Data, vol. 8, no. 2, pp. 302–317, 2022.
- Y. Wang, A. Davidson, Y. Pan, Y. Wu, A. Riffel, and J. D. Owens, "Gunrock: A high-performance graph processing library on the gpu," in Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2016, pp. 1–12
- Jeremy D. Wendt, Richard Field, Cynthia Phillips, Arvind Prasadan, Tegan Wilson, Sucheta Soundarajan, and Sanjukta Bhowmick. 2023. Partitioning Commnuica- tion Streams Into Graph Snapshots. IEEE Transactions on Network Science and Engineering 10, 2 (2023), 809–826. https://doi.org/10.1109/TNSE.2022.3223614

Acknowledgment

This work was partially supported by the NSF projects SANDY (Award # OAC-1725755) and CANDY (Award # OAC-2104115, 2104076, 2104078). We extend our gratitude to Jeremy Wendt, Richard Field, Cynthia Philips, Arvind Prasadan and Sucheta Soundarajan for their guidance on generating synthetic dynamic graphs.