Modeling Single-Source Shortest Path Algorithm Dynamics to Control Performance And Power Tradeoffs

Sara Karamati, Jeffrey Young (advisor), Richard Vuduc (advisor)
Georgia Institute of Technology, Atlanta, Georgia

Key Idea

Design dynamical system models that can automatically tune an algorithm for optimal performance and across diverse hardware combinations.

Use Case

GPU-based parallel single-source shortest path (SSSP) from the Gunrock [1] analytics library.

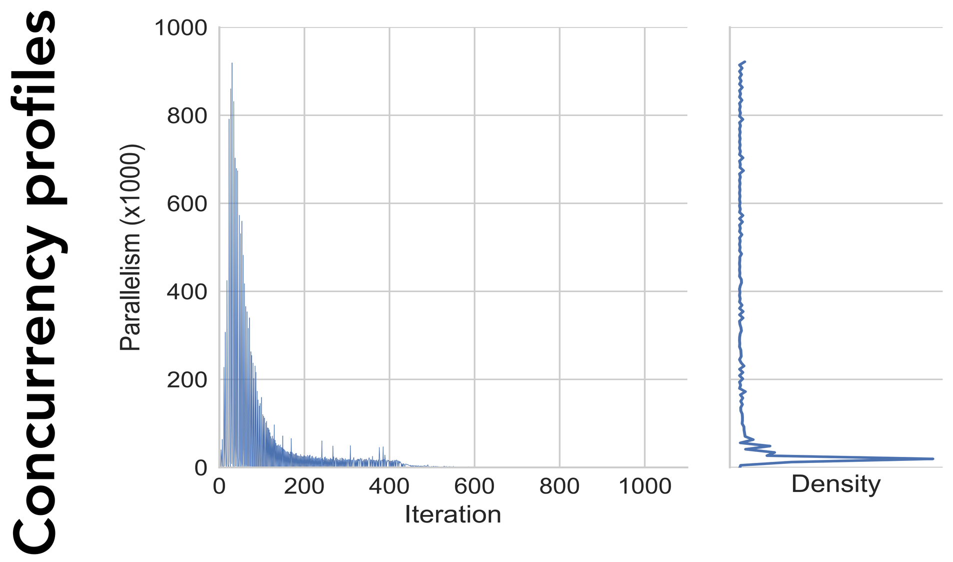

The base line Near+Far SSSP [2]

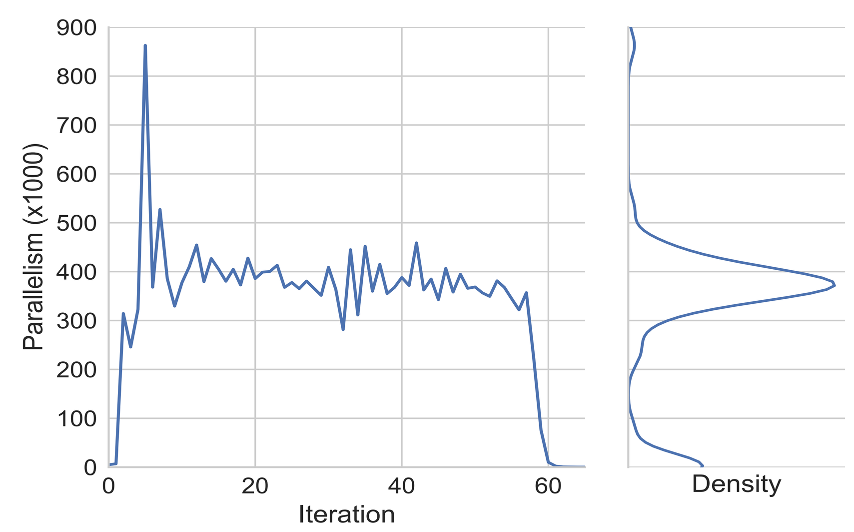

Our proposed SSSP algorithm [3]

Proposed controller keeps parallelism consistent across iterations

A Self-Tuning Algorithm

- ✓ Develop a mathematical model that uses online learning techniques to predict the available parallelism in each iteration based on the tunable delta parameter.
- ✓ Use this model to tune the available parallelism dynamically to meet a given target, thereby improving the average available parallelism while reducing its variability.

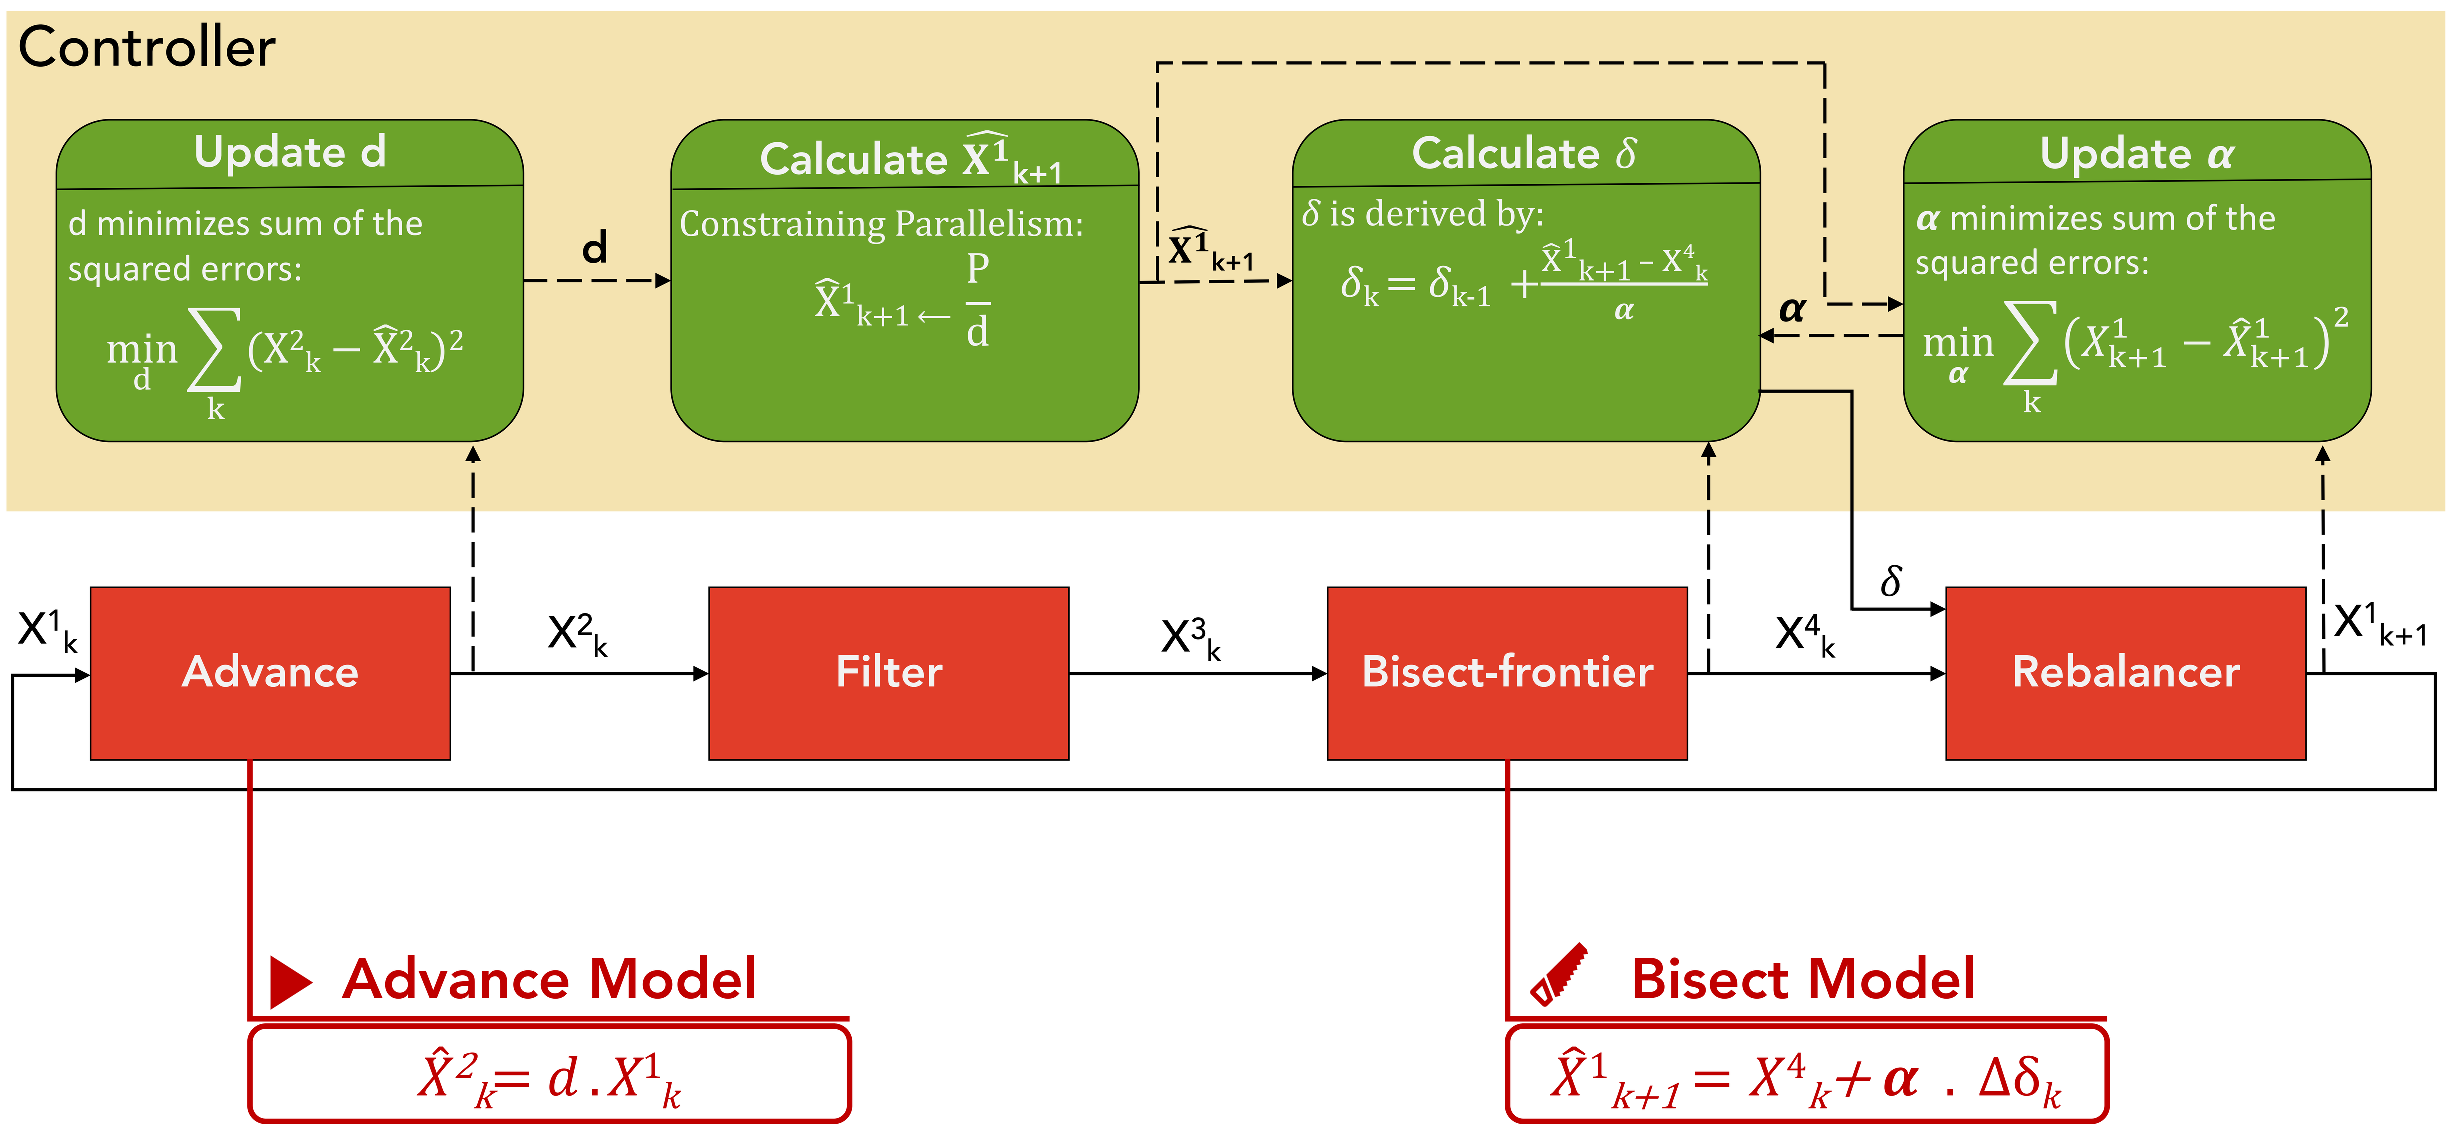

Performance and Power

Cal road network [4]

- ✓ Controlled parallelism improves the performance and power consumption over the baseline.
- ✓ Most of the self-tuning points are both faster and more energy-efficient than the baseline algorithm.

Wikipedia-20051105 network [5]

- ✓ Algorithmic knobs yield a 50% speedup with a 25% power savings.

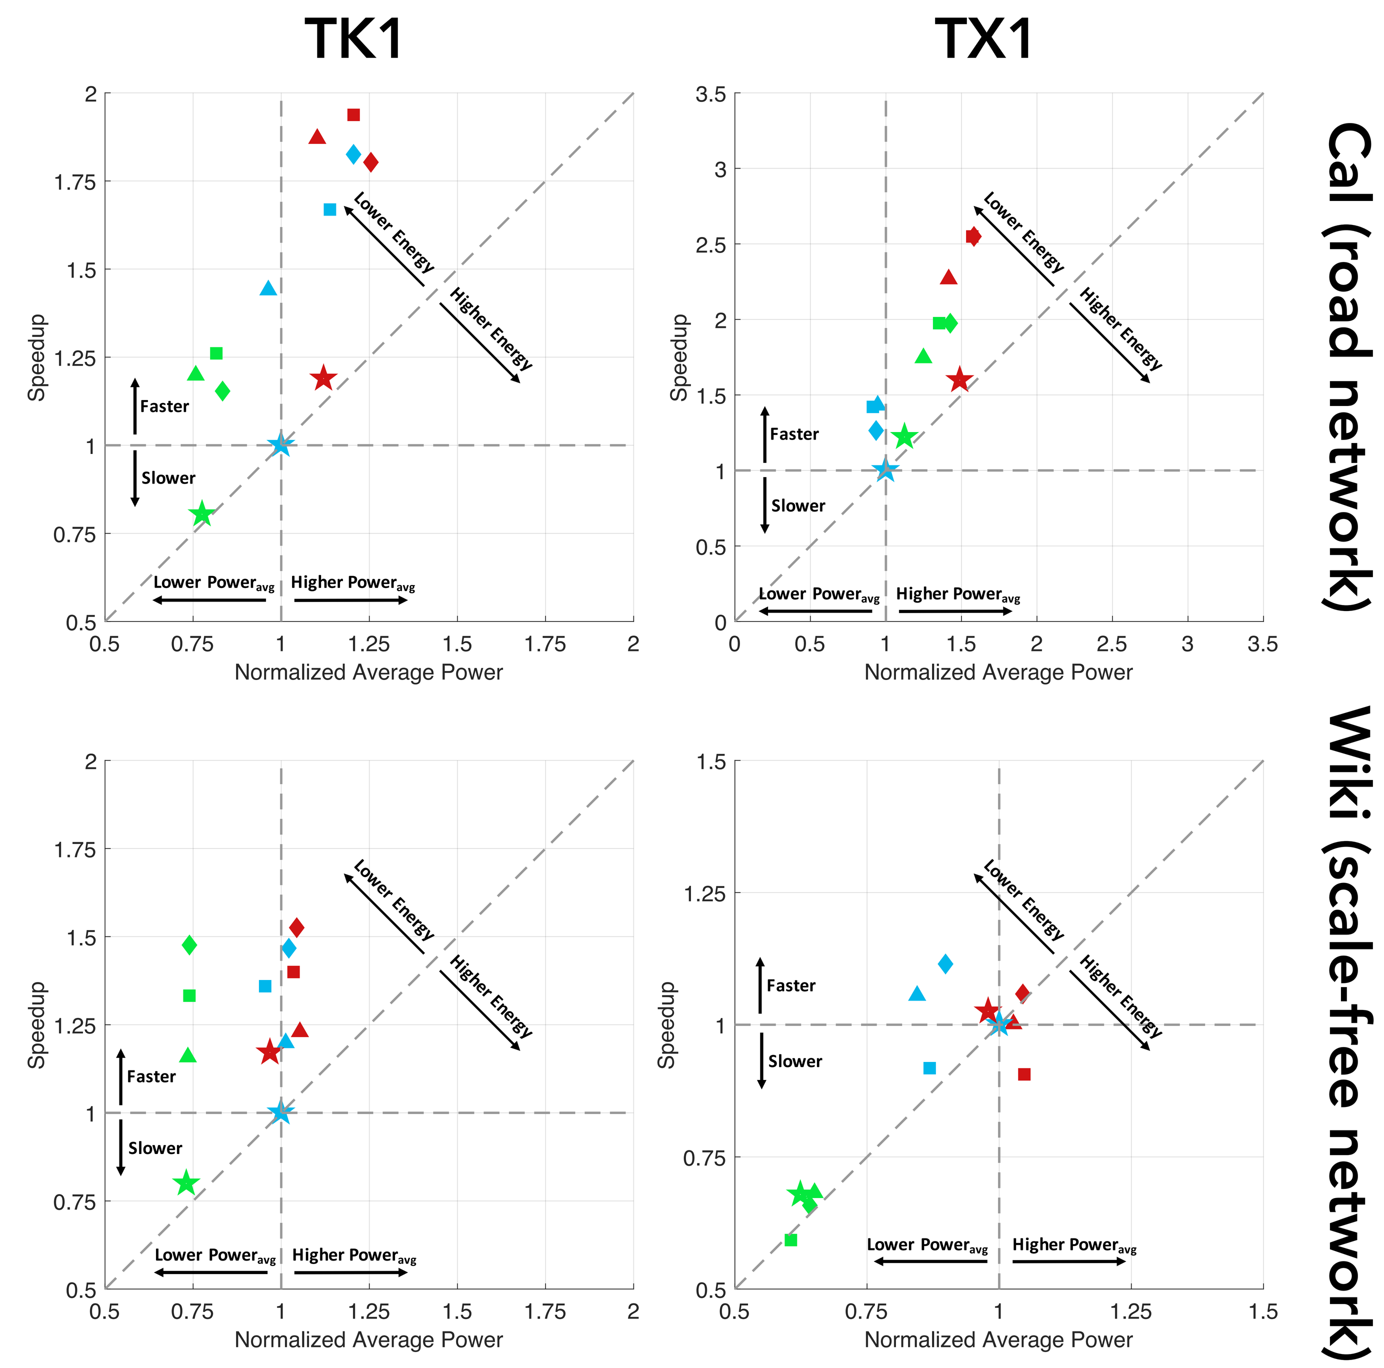

Efficacy of parallelism control

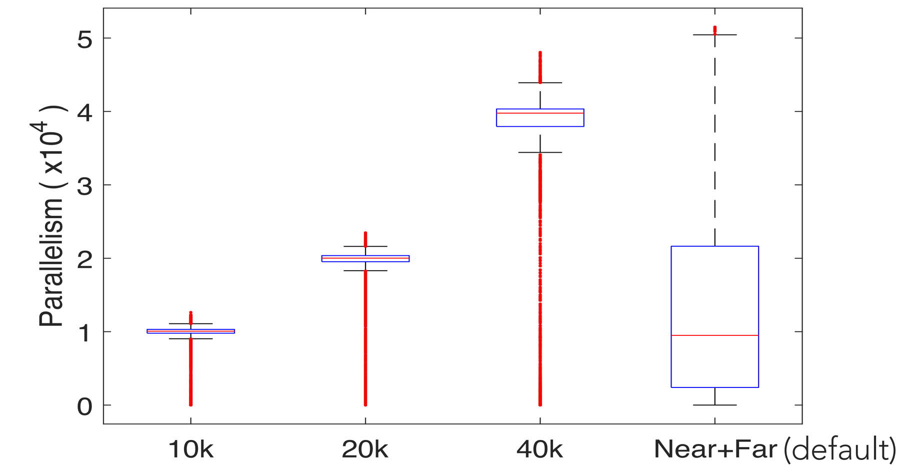

Proposed technique limits the distribution of available parallelism for the Cal road network

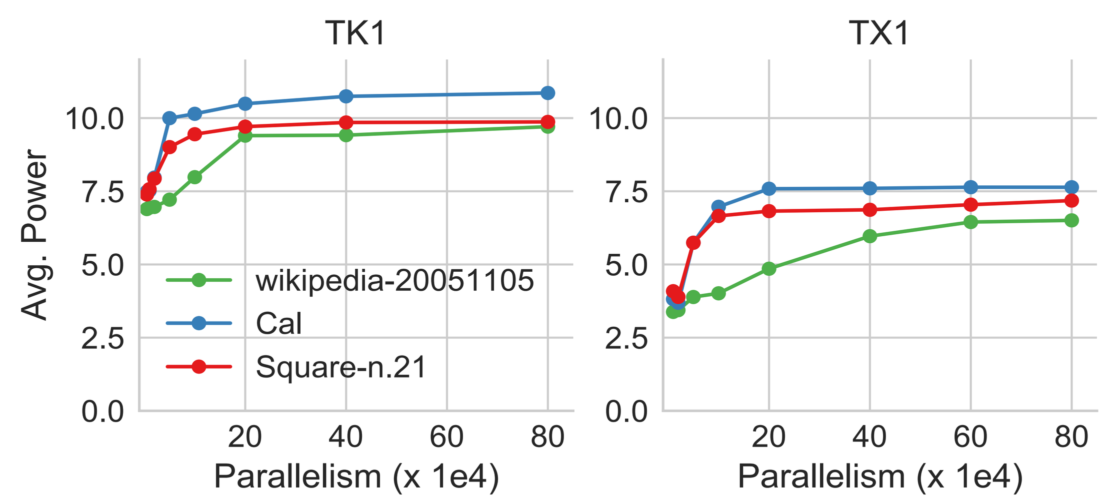

Power shows a sublinear dependency on the parallelism

Conclusion

- ✓ It is possible to design an automatic controller scheme to manage irregular parallelism.
- ✓ The controller may be used to improve energy-efficiency or limit power consumption as part of increasing performance or managing a power-performance tradeoff.
- ✓ Designing such a technique at the level of the algorithm complements power-performance management at other levels, such as at the hardware level through DVFS mechanisms.

1.Wang, Yangzihao, et al. "Gunrock: A high-performance graph processing library on the GPU." ACM SIGPLAN Notices. Vol. 51. No. 8. ACM, 2016.
2.Davidson, Andrew Alan, et al. "Work-efficient parallel GPU methods for single-source shortest paths." (2014).
3.Karamati, Sara, Jeffrey Young, and Richard Vuduc. "An Energy-Efficient Single-Source Shortest Path Algorithm." 2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS). IEEE, 2018.
4.C. Demetrescu, A. V. Goldberg, and D. S. Johnson, The Shortest Path Problem: Ninth DIMACS Implementation Challenge. American Mathematical Soc., vol. 74.
5.T. A. Davis and Y. Hu, "The university of florida sparse matrix collection," ACM Trans. Math. Softw., vol. 38, no. 1, pp. 1:1–1:25, Dec. 2011. [Online]. Available: http://doi.acm.org/10.1145/2049662.2049663