Scalable Construction of Spectrally Near-Optimal Networks via Reinforcement Learning

Junsung Hwang,[1] Ji Liu [1]

1. Stony Brook University, Stony Brook, NY, USA

The algebraic connectivity of a simple graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with $N = |\mathcal{V}|$ nodes and $M = |\mathcal{E}|$ edges, fundamentally governs the robustness of networked systems, controlling synchronization speeds in power grids, consensus convergence in distributed computing, and information diffusion rates [1]. While maximizing λ_2 is NP-hard [2], greedy spectral heuristics, such as maximizing Effective Resistance (ER) or Fiedler Vector (FV) separation, can produce high-quality topologies. However, these methods require expensive eigendecompositions or matrix inversions ($O(N^3)$), resulting in a prohibitive total complexity of $O(MN^3)$ or $O(N^5)$ for dense graphs. Conversely, scalable $O(N^2)$ constructive models like the Watts-Strogatz Small World (SW) often fail to avoid spectral bottlenecks, yielding suboptimal robustness.

We introduce a Two-Phase Hybrid Framework that delivers high algebraic connectivity with $O(N^2)$ inference, bridging the gap between spectral quality and computational scalability.

Phase I (Deterministic Skeleton): We propose a deterministic strategy that approximates high spectral quality via local geometric proxies. By prioritizing edges with maximal chordal distance and disjoint neighborhoods, we facilitate isotropic expansion, promoting robust, monotonic connectivity scaling across arbitrary edge densities.

Phase II (RL Refinement): To surpass deterministic limits, we employ a Graph Convolutional Network (GCN) agent trained via Reinforcement Learning to add residual edges. Crucially, we introduce a Top-K Guided Policy: the deterministic heuristic from Phase I prunes the search space to the top K candidates, allowing the RL agent to focus on learning subtle topological optimizations within a high-quality subspace. This imparts guided flexibility to the pipeline without sacrificing scalability.

We benchmark our method across validation scales $N \in \{32, 64, 128, 256\}$, where greedy spectral baselines remain tractable for direct comparison, since ER and FV heuristics recompute a Laplacian eigensolver at each greedy edge-addition step. Results show that our $O(N^2)$ hybrid pipeline achieves 95–97% of the greedy ER baseline's λ_2 while consistently outperforming Small-World graphs by 10–20%. We further evaluate scalability up to $N = 2{,}048$, observing quadratic runtime scaling and sustained λ_2 gains of approximately 15% over Small-World baselines.

[1] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.
[2] D. Mosk-Aoyama, "Maximum algebraic connectivity augmentation is NP-hard," *Operations Research Letters*, vol. 36, no. 6, pp. 677–679, 2008.
[3] A. Ghosh and S. Boyd, "Growing well-connected graphs," *Proceedings of the 45th IEEE Conference on Decision and Control*, pp. 6605–6611, 2006.

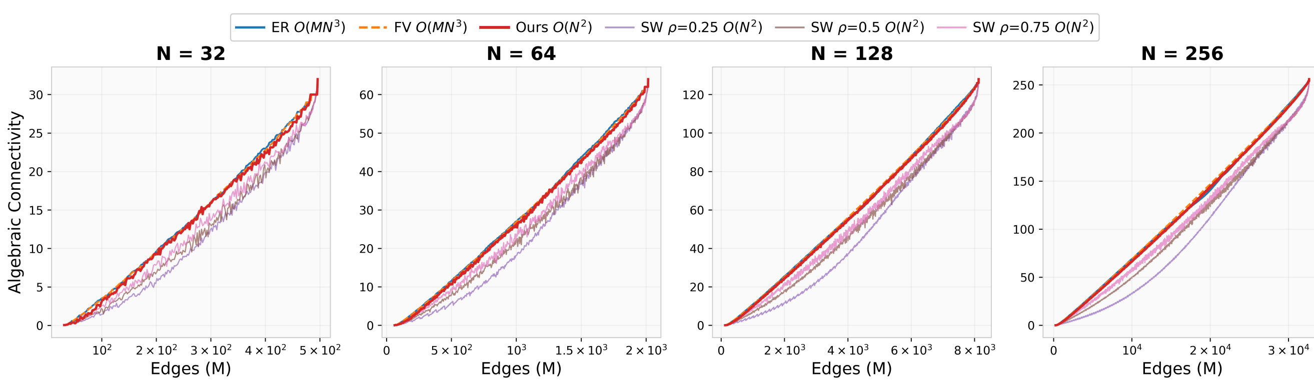

Figure 1: Algebraic connectivity λ_2 versus edge count M for graphs of size $N \in \{32, 64, 128, 256\}$. We compare greedy methods (Effective Resistance and Fiedler Vector, both $O(MN^3)$) that optimize λ_2 at each edge addition against our $O(N^2)$ pipeline and Small-World baselines with rewiring probability $\rho \in \{0.25, 0.5, 0.75\}$.