Triangle Counting with Local Edge Differential Privacy

Friday, July 14, 2023 12:10 PM (20 minutes)

Talya Eden, Sofya Raskhodnikova, Quanquan C. Liu and Adam Smith

Abstract: Many deployments of differential privacy in industry are in the local model, where each party releases its private information via a differentially private randomizer. We study triangle counting in the noninteractive and interactive local model with edge differential privacy (that, intuitively, requires that the outputs of the algorithm on graphs that differ in one edge be indistinguishable). In this model, each party's local view consists of the adjacency list of one vertex.

In the noninteractive model, we prove that additive $\Omega(n^2)$ error is necessary, where n is the number of nodes. This lower bound is our main technical contribution. It uses a reconstruction attack with a new class of linear queries and a novel mix-and-match strategy of running the local randomizers with different completions of their adjacency lists. It matches the additive error of the algorithm based on Randomized Response, proposed by Imola, Murakami and Chaudhuri (USENIX2021) and analyzed by Imola, Murakami and Chaudhuri (CCS2022) for constant ε . We use a different postprocessing of Randomized Response and provide tight bounds on the variance of the resulting algorithm. In the interactive setting, we prove a lower bound of $\Omega(n^{3/2})$ on the additive error. Previously, no hardness results were known for interactive, edge-private algorithms in the local model, except for those that follow trivially from the results for the central model. Our work significantly improves on the state of the art in differentially private graph analysis in the local model.

Presenter: EDEN, Talya

Session Classification: Track A-3