Locally Private k-Means in One Round*

Alisa Chang Badih Ghazi Ravi Kumar Pasin Manurangsi
Google, Mountain View, CA.
{alisac, pasin}@google.com, {badihghazi, ravi.k53}@gmail.com

Summary of Results. We provide an approximation algorithm for k-means clustering in the *one-round* (aka *non-interactive*) local model of differential privacy (DP). Our algorithm achieves an approximation ratio arbitrarily close to the best *non private* approximation algorithm, improving upon previously known algorithms that only guarantee large (constant) approximation ratios. Furthermore, ours is the first constant-factor approximation algorithm for k-means that requires only *one* round of communication in the local DP model, positively resolving an open question of Stemmer (SODA, 2020). Our algorithmic framework is quite flexible; we demonstrate this by showing that it also yields a similar near-optimal approximation algorithm in the (one-round) shuffle DP model.

*The full version of this paper will appear in ICML 2021.

2nd Symposium on Foundations of Responsible Computing (FORC), June 9-11, 2021