Abstract

Miller, Peng, Vladu and Xu, (Improved Parallel Algorithms for Spanners and Hopsets), SPAA'15, devised (implicitly) a distributed algorithm in the CONGEST model, that given a parameter $k = 1, 2, \ldots$, constructs an $O(k)$-spanner of an input unweighted n-vertex graph with $O(n^{1+1/k})$ expected edges in $O(k)$ rounds of communication. In this paper we improve the result of Miller et al., by showing a k-round distributed algorithm in the same model, that constructs a $(2k-1)$-spanner with $O(n^{1+1/k}/\epsilon)$ edges, with probability $1-\epsilon$, for any $\epsilon > 0$. Moreover, when $k = \omega(\log n)$, our algorithm produces (still in k rounds) *ultra-sparse* spanners, i.e., spanners of size $n(1 + o(1))$, with probability $1 - o(1)$. To our knowledge, this is the first distributed algorithm in the CONGEST or in the PRAM models that constructs spanners or skeletons (i.e., connected spanning subgraphs) that sparse. Our algorithm can also be implemented in linear time in the standard centralized model, and for large k, it provides spanners that are sparser than any other spanner given by a known (near-)linear time algorithm. We also devise improved bounds (and algorithms realizing these bounds) for $(1 + \epsilon, \beta)$-spanners and emulators. In particular, we show that for any unweighted n-vertex graph and any $\epsilon > 0$, there exists a $(1 + \epsilon, (\frac{\log\log n}{\epsilon})^{\log\log n})$-emulator with $O(n)$ edges. All previous constructions of $(1 + \epsilon, \beta)$-spanners and emulators employ a super-linear number of edges, for all choices of parameters. Finally, we provide some applications of our results to approximate shortest paths' computation in unweighted graphs.