Abstract

The Tree Sparsity problem is defined as follows: given a node-weighted tree of size n and an integer k, output a rooted subtree of size k with maximum weight. The best known algorithm solves this problem in time $O(kn)$, i.e., quadratic in the size of the input tree for $k = \Theta(n)$. In this work, we design $(1+\varepsilon)$-approximation algorithms for the Tree Sparsity problem that run in nearly-linear time. Unlike prior algorithms for this problem, our results offer single criterion approximations, i.e., they do not increase the sparsity of the output solution, and work for arbitrary trees (not only balanced trees). We also provide further algorithms for this problem with different runtime vs approximation trade-offs. Finally, we show that if the exact version of the Tree Sparsity problem can be solved in strongly subquadratic time, then the $(\min, +)$ convolution problem can be solved in strongly subquadratic time as well. The latter is a well-studied problem for which no strongly subquadratic time algorithm is known.