Abstract

In this paper we consider the decremental single-source shortest paths (SSSP) problem, where given a graph G and a source node s the goal is to maintain shortest paths between s and all other nodes in G under a sequence of online adversarial edge deletions. (Our algorithm can also be modified to work in the incremental setting, where the graph is initially empty and subject to a sequence of online adversarial edge insertions.) In their seminal work, Even and Shiloach [JACM 1981] presented an exact solution to the problem with only $O(mn)$ total update time over all edge deletions. Later papers presented conditional lower bounds showing that $O(mn)$ is optimal up to log factors. In SODA 2011, Bernstein and Roditty showed how to bypass these lower bounds and improve upon the Even and Shiloach $O(mn)$ total update time bound by allowing a $(1+\epsilon)$ approximation. This triggered a series of new results, culminating in a recent breakthrough of Henzinger, Krinninger and Nanongkai [FOCS 14], who presented a $(1+\epsilon)$-approximate algorithm whose total update time is near linear: $O(m^{1+o(1)})$. However, every single one of these improvements over the Even-Shiloach algorithm was randomized and assumed a non-adaptive adversary. This additional assumption meant that the algorithms were not suitable for certain settings and could not be used as a black box data structure. Very recently Bernstein and Chechik presented in STOC 2016 the first *deterministic* improvement over Even and Shiloach, that did not rely on randomization or assumptions about the adversary: in an undirected unweighted graph the algorithm maintains $(1+\epsilon)$-approximate distances and has total update time $\tilde{O}(n^2)$. In this paper, we present a new deterministic algorithm for the problem with total update time $\tilde{O}(n^{1.25}\sqrt{m}) = \tilde{O}(mn^{3/4})$: it returns a $(1+\epsilon)$ approximation, and is limited to undirected unweighted graphs. Although this result is still far from matching the randomized near-linear total update time, it presents important progress towards that direction, because unlike the STOC 2016 $\tilde{O}(n^2)$ algorithm it beats the Even and Shiloach $O(mn)$ bound for all graphs, not just sufficiently dense ones. In particular, the $\tilde{O}(n^2)$ algorithm relied entirely on a new sparsification technique, and so could not hope to yield an improvement for sparse graphs. We present the first deterministic improvement for sparse graphs by significantly extending some of the ideas from the $\tilde{O}(n^2)$ algorithm and combining them with the hop-set technique used in several earlier dynamic shortest path papers.