Abstract

Spanners, emulators, and approximate distance oracles can be viewed as *lossy* compression schemes that represent an unweighted graph metric in small space, say $\tilde{O}(n^{1+\delta})$ bits. There is an inherent tradeoff between the sparsity parameter δ and the *stretch function* f of the compression scheme, but the qualitative nature of this tradeoff has remained a persistent open problem. It has been known for some time that when $\delta \geq 1/3$ there are schemes with constant *additive* stretch (distance d is stretched to at most $f(d) = d+O(1)$), and recent results of Abboud and Bodwin show that when $\delta < 1/3$ there are no such schemes. Thus, to get practically efficient graph compression with $\delta \to 0$ we must pay super-constant additive stretch, but exactly how much do we have to pay? In this paper we show that the lower bound of Abboud and Bodwin is just the first step in a *hierarchy* of lower bounds that characterize the asymptotic behavior of the optimal stretch function f for sparsity parameter $\delta \in (0, 1/3)$. Specifically, for any integer $k \geq 2$, any compression scheme with size $O(n^{1+\frac{1}{2^k-1}-\epsilon})$ has a *sublinear additive stretch* function f:

$$f(d) = d + \Omega(d^{1-\frac{1}{k}}).$$

This lower bound matches Thorup and Zwick's (2006) construction of sublinear additive *emulators*. It also shows that Elkin and Peleg's $(1+\epsilon, \beta)$-spanners have an essentially optimal tradeoff between δ, ϵ, and β, and that the sublinear additive spanners of Pettie (2009) and Chechik (2013) are not too far from optimal. To complement these lower bounds we present a new construction of $(1 + \epsilon, O(k/\epsilon)^{k-1})$-spanners with size $O((k/\epsilon)^{h_k} kn^{1+\frac{1}{2^{k+1}-1}})$, where $h_k < 3/4$. This size bound improves on the spanners of Elkin and Peleg (2004), Thorup and Zwick (2006), and Pettie (2009). According to our lower bounds neither the size nor stretch function can be substantially improved.