If we have found a "good" clustering $C$ of data set $X$, can we prove that $C$ is not far from the (unknown) best clustering $C^*$ of this data set? Perhaps surprisingly, the answer to this question is sometimes yes. We can show bounds on the distance $d(C, C^*)$ for two clustering cost functions: the Normalized Cut and the squared distance cost of K-means clustering. These bounds exist in the case when the data $X$ admits a "good" clustering for the given cost.