Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institute of Theoretical Computer Science
Mohsen Ghaffari, Angelika Steger, Emo Welzl, Peter Widmayer

Solution 1: MIS via Network Decomposition

Given: A (C_n, D_n)-network decomposition consisting of blocks $G_1, \dots, G_{C_n}$. (By "given" we mean that every node knows the index of the block that it belongs to, as well as the index of its cluster within that block.)

First we argue that we can compute a maximal independent set of G_1 in time $O(D_n)$. Indeed, in order to compute a maximal independent set of one cluster X_j, it suffices that every node in X_j gets to know what the vertices and edges in X_j are; and since any two vertices in X_j have distance at most D_n, this can be done in $O(D_n)$ rounds. Recall that the clusters within G_1 are pairwise non-adjacent: Hence, if we compute a maximal independent set of each cluster X_j "in parallel", then we obtain a maximal independent set of G_1.

Now assume that we have already computed an MIS $M_{i-1} \subseteq V$ for $G_1 \cup \cdots \cup G_{i-1}$, and we want to observe that we can extend this to an MIS M_i for $G_1 \cup \cdots \cup G_i$, in $O(D_n)$ additional results. This is the same argument as before: After $D_n + 1$ rounds the vertices in every cluster X_j of G_i know what are the vertices and edges within X_j as well as what are the edges between X_j and M_{i-1}. Based on this information they can compute a maximal subset $M_{i,j} \subseteq X_j$ such that $M_{i-1} \cup M_{i,j}$ is still an independent set; and $M_i := M_{i-1} \cup \bigcup_j M_{i,j}$ is an independent set of $G_1 \cup \cdots \cup G_i$.

Repeating this for $i = 2, \dots, C_n$ gives the bound $O(C_n D_n)$.

Solution 2: Near-Optimality of Theorem 8.27

Note: There was a mistake in the question that reversed the roles of C *and* D. *We actually prove that* (C_n, D_n)*-network decompositions with* $C_n = o(\log n)$ *and* $D_n = o\left(\frac{\log n}{\log \log n}\right)$ *do not exist.*

We prove the question for *strong* network decompositions. (The proof for weak network decompositions require just one additional step; see the remark below.)

Let the given graphs be called $G(n)$ and assume the opposite, i. e., assume there are strong (C_n, D_n)-network decompositions of $G(n)$ with $C_n \in o(\mathrm{chrom}(G(n)))$ and $D_n \in o(\mathrm{girth}(G(n)))$. Then in particular there is $n_0 \in \mathbf{N}$ such that for all $n \geq n_0$,

$$C_n < \frac{1}{2}\,\mathrm{chrom}(G(n)) \quad \text{and} \quad D_n < \frac{1}{2}\,\mathrm{girth}(G(n)) - 1.$$

Fix $n \geq n_0$. For the sake of finding a contradiction we want to exhibit a coloring of $G(n)$ that uses less than $\mathrm{chrom}(G)$ colors.

Consider a cluster X in $G(n)$. We claim that X is a forest. Indeed, if there is a cycle in X, choose a cycle of minimum length, and consider two vertices u, v that lie both on the cycle, and such that u, v have maximum distance along the cycle (they are "opposite" vertices on the cycle). The vertices u and v break the cycle into two paths from u to v, both of length at least $\frac{1}{2}\,\mathrm{girth}(G(n))-1$. By definition of strong network decompositions there must be another path in X from u to v of length at most $D_n < \frac{1}{2}\,\mathrm{girth}(G(n)) - 1$. This path will necessarily create a shortcut in our cycle: a contradiction to minimality.

We have shown that every cluster is a forest, hence it can be colored with 2 colors. Since the clusters within each block are non-adjacent, the whole block can be colored with 2 colors and hence G can be colored with $2C_n < \mathrm{chrom}(G)$ colors, a contradiction.

Remark. How to prove that also *weak* (C_n, D_n)-network decompositions do not exist? At first sight the proof above breaks down because the shortcut in the cycle might not lie within X. But this issue can easily be circumvented like this: Instead of proving directly that the cluster X is a forest, we take any vertex $v \in X$ and then we prove that the set $\{u \in V : \mathrm{dist}_G(u, v) \leq D_n\}$ is a forest. This set is a superset of X, so we can deduce that also X is a forest.

Solution 3: Diameter Orderings

Consider a (C, D)-network composition as in theorem 8.27, where $C = O(\log n)$ and $D = O(\log n)$, with blocks $G_1, \ldots, G_C$. Let f be any ordering with the property that, for all $i < j$, everything in G_i is smaller than anything in G_j.

In order to show that f is an $O(\log^2 n)$-diameter ordering, let $v_1, \ldots, v_p$ be a path that is monotone with respect to f. We want to show that $\mathrm{dist}(v_1, v_p) \leq CD$. By construction of f the path can be broken up into subsequences (some of which may be empty)

$$\begin{aligned} v_1, \ldots, v_{i_1} &\in G_1, \\ v_{i_1+1}, \ldots, v_{i_2} &\in G_2, \\ &\ldots, \\ v_{i_{C-1}+1}, \ldots, v_{i_C} = v_p &\in G_C. \end{aligned}$$

Since the clusters of each G_i are pairwise non-adjacent, each subsequence is entirely contained in one cluster. Hence

$$\mathrm{dist}(v_1, v_{i_1}) \leq D, \quad \mathrm{dist}(v_{i_1+1}, v_{i_2}) \leq D, \quad \ldots, \quad \mathrm{dist}(v_{i_{C-1}+1}, v_{i_C}) \leq D;$$

and

$$\mathrm{dist}(v_1, v_p) \leq \mathrm{dist}(v_1, v_{i_1}) + \cdots + \mathrm{dist}(v_{i_{C-1}+1}, v_{i_C}) \leq CD \in O(\log^2 n).$$

Solution 4: Ruling Sets

By definition, S is a $(2, O(\log n))$-ruling set for $W = V$ if

(i) every two vertices in S have distance at least 2 in G, and

(ii) for every vertex $v \in V \setminus S$ there is $u \in S$ such that $\mathrm{dist}_G(v, u) \leq O(\log n)$.

Condition (i) is immediate from the construction; two adjacent vertices cannot be both strict local minima.

We now prove that condition (ii) holds with high probability. Let $v \in V \setminus S$ and consider, for some $k \in \mathbf{N}$, the "k-neighborhood"

$$N^k(v) := \{u \in V : \mathrm{dist}(v, u) \leq k\}.$$

Furthermore let $x_k \in N^k(v)$ denote the smallest (w.r.t. f) vertex in the k-neighborhood.

Observation: If $x_k \in N^{k-1}(v)$ holds, then x_k is a local minimum. (You may want to draw a picture in order to convince yourself of this statement.)

Motivated by the observation we are interested in the event that, say, for some $k \in \{1, 2, \ldots, 4 \log_n\}$, the event $x_k \in N^{k-1}(v)$ takes place. Since every vertex in $N^k(v)$ is equally likely to be the smallest vertex x_k, we have

$$\Pr\left[x_k \in N^{k-1}(v)\right] = \frac{|N^{k-1}(v)|}{|N^k(v)|}.$$

How do we bound this probability from below? In general we cannot; it can be a nearly arbitrary ratio between 0 and 1. But assume that *there are more than* $\log n$ *distinct values for* $k \in \{1, 2, \ldots, 4 \log_n\}$ *with* $\frac{|N^{k-1}(v)|}{|N^k(v)|} < \frac{1}{2}$. Then $\left|N^{4 \log n}(v)\right| > 2^{\log n} = n$, an impossibility. Thus we have shown that for at least $3 \log n$ distinct values of $k \in \{1, 2, \ldots, 4 \log n\}$ we have

$$\frac{|N^{k-1}(v)|}{|N^k(v)|} \geq \frac{1}{2}.$$

We calculate

$$\begin{aligned}
&\Pr\left[\text{ the set } N^{4 \log n}(v) \text{ contains a local minimum }\right] \\
&\geq \Pr\left[\bigcup_{k=1}^{4 \log n} x_k \in N^{k-1}(v)\right] = 1 - \Pr\left[\bigcap_{k=1}^{4 \log n} x_k \notin N^{k-1}(v)\right] \\
&\overset{(*)}{=} 1 - \prod_{k=1}^{4 \log n} \Pr\left[x_k \notin N^{k-1}(v)\right] \geq 1 - \left(\frac{1}{2}\right)^{3 \log n} \cdot 1^{\log n} = 1 - \frac{1}{n^3}.
\end{aligned}$$

We skip the proof for $(*)$; a very brief and very informal justification: For the probability that x_k lies on the boundary of $N^k(v)$ it does not matter how the numbers inside $N^{k-1}(v)$ are assigned to the individual vertices.

The previous calculation was valid for every vertex v. Now we take a union bound over all vertices:

$$\begin{aligned}
&\Pr\left[\forall v \in V\ \exists u \in S : \mathrm{dist}_G(v, u) \leq 4 \log n\right] \\
&\geq 1 - \Pr\left[\bigcup_{v \in V} \{\nexists u \in S : \mathrm{dist}_G(v, u) \leq 4 \log n\ \}\right] \\
&\geq 1 - n \cdot \frac{1}{n^3} = 1 - \frac{1}{n^2}.
\end{aligned}$$