

The Query Complexity of Finding a Hidden Permutation

KURT MEHLHORN

Max-Planck-Institut für Informatik

We study the query complexity of determining a hidden permutation. More specifically, we study the problem of learning a secret (z, π) consisting of a binary string z of length n and a permutation π of $[n]$. The secret must be unveiled by asking queries $x \in \{0, 1\}^n$, and for each query asked, we are returned the score $f_{z, \pi}(x)$ defined as

$$f_{z, \pi}(x) := \max\{i \in [0..n] \mid \forall j \leq i : z_{\pi(j)} = x_{\pi(j)}\};$$

i.e., the length of the longest common prefix of x and z with respect to π . The goal is to minimize the number of queries asked. We prove matching upper and lower bounds for the deterministic and randomized query complexity of $\Theta(n \log n)$ and $\Theta(n \log \log n)$, respectively.

Joint work with Peyman Afshani, Manindra Agrawal, Benjamin Doerr, Carola Doerr, and Kasper Green Larsen