Trace reconstruction for the deletion channel

Yuval Peres [1],

[1] Microsoft Research

In the trace reconstruction problem, an unknown string x of n bits is observed through the deletion channel, which deletes each bit with some constant probability q, yielding a contracted string. How many independent outputs (traces) of the deletion channel are needed to reconstruct x with high probability?

The best lower bound known is linear in n. Until 2016, the best upper bound was exponential in the square root of n. In earlier work with F. Nazarov (STOC 2017), we improved the square root to a cube root using statistics of individual output bits and some inequalities for Littlewood polynomials on the unit circle. This bound is sharp for reconstruction algorithms that only use this statistical information. (Similar results were obtained independently and concurrently by De, O'Donnell and Servedio). If the string x is random, we can show a subpolynomial number of traces suffices by comparison to a random walk. (Joint works with Alex Zhai, FOCS 2017 and with Nina Holden Robin Pemantle, preprint (2017).)