

Tiny Families of Functions with Random Properties: A Quality-Size Trade-off for Hashing

Oded Goldreich
A. Wigderson

Abstract

We present three explicit constructions of hash functions, which exhibit a trade-off between the size of the family (and hence the number of random bits needed to generate a member of the family), and the quality (or error parameter) of the pseudo-random property it achieves. Unlike previous constructions, most notably universal hashing, the size of our families is essentially independent of the size of the domain on which the functions operate.

The first construction is for the *mixing* property – mapping a proportional part of any subset of the domain to any other subset. The other two are for the *extraction* property – mapping any subset of the domain almost uniformly into a range smaller than it. The second and third constructions handle (respectively) the extreme situations when the range is very large or very small.

We provide lower bounds showing that our constructions are nearly optimal, and mention some applications of the new constructions.