Randomized Algorithms 2020-1
Lecture 11

Cuckoo Hashing *

Moni Naor

1 Recap

We reviewed the applications of the local lemma and to showing that CNFs with few appearance of each variable are satisfiable.

2 Hash Tables

We mentioned the important open problem of showing that randomness is essential to achieve efficient ($O(1)$) dictionaries.

In general, the type of error probability we have in the is area is $1/n^c$ (and not exponentially small).

We went over Cuckoo Hashing and its analysis. We discussed what is needed to detect failure. In particular mentioned that a good cycle detection mechanism is useful.

Question: suggest a low memory cycle detection.

The probability of failure of cuckoo hashing after n insertions is $1/n$. We talked about the idea of a stash to reduce the probability using a stash. This reduces the failure probability to $1/n^{c+1}$ where c is the size of the stash.

Question: We saw the proof that if there are no large components, then the probability of failure (= a component with 2 cycles). Extend the argument to show that the probability one needs to put c elements into the stash is rough $1/n^{c+1}$.

References

[1] J. L. Carter and M. N. Wegman, Universal classes of hash functions, J. Comput. Syst. Sci. 18 (1979) 143–154.

*These notes summarize the material covered in class, usually skipping proofs, details, examples and so forth, and possibly adding some remarks, or pointers. In the interest of brevity, most references and credits were omitted.

http://www.cs.princeton.edu/courses/archive/fall09/cos521/Handouts/universalclasses.pdf

[2] Kirsch, Mitzenmacher, Wieder, *More robust hashing: Cuckoo hashing with a stash*, ESA 2008.

[3] Rasmus Pagh and Flemming Friche Rodler: *Cuckoo hashing*. J. Algorithms 51(2): 122-144 (2004).

[4] R. Panigrahy, *Efficient hashing with lookups in two memory accesses*, SODA 2005