HASH FILTER: MAKING BLOOM FILTER EVEN MORE COMPACT AND DELETABLE

Bin Fan, David G. Andersen (Carnegie Mellon University), Michael Kaminsky (Intel Labs Pittsburgh)

BLOOM FILTER: APPROXIMATE SET-MEMBERSHIP TEST

- **Answers questions like "is foo present in this set?"**
- **Returns "no" with 100% certianty, "yes" with a false positive rate e (tunable parameter)**

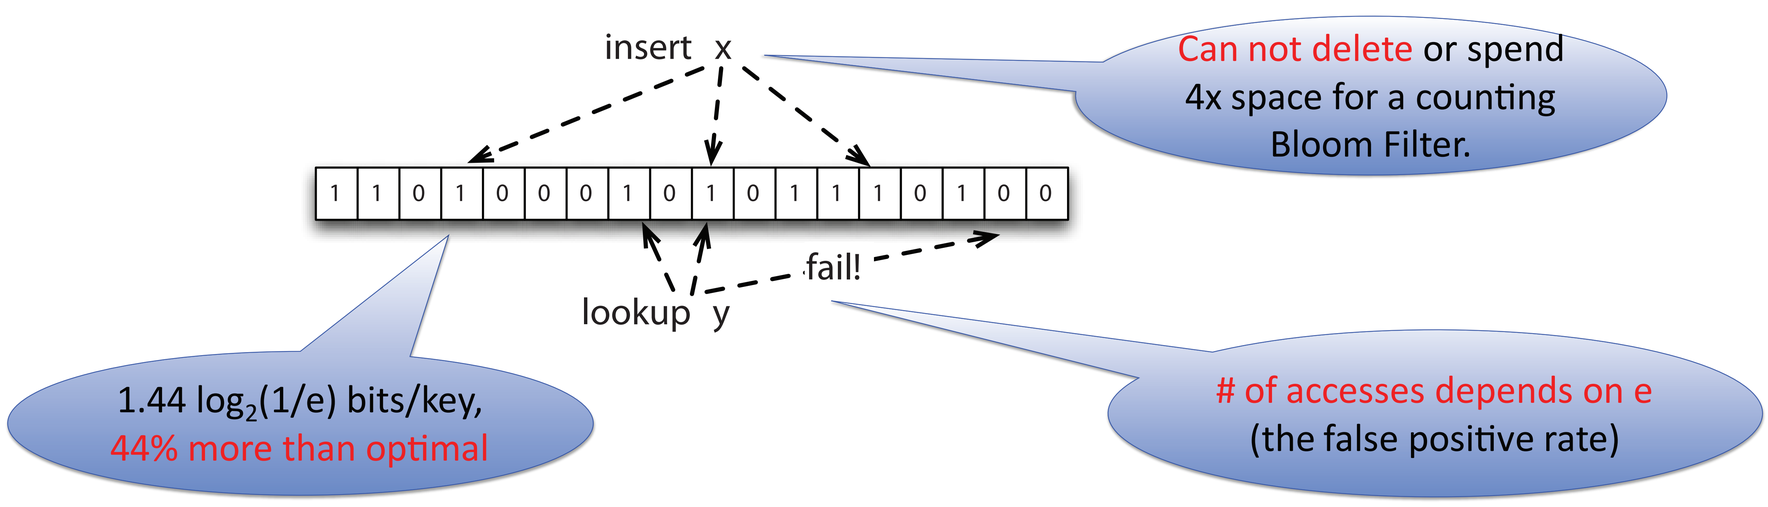

HASH FILTER IN A NUTSHELL

- **Essentially a hash table filled with signatures!**

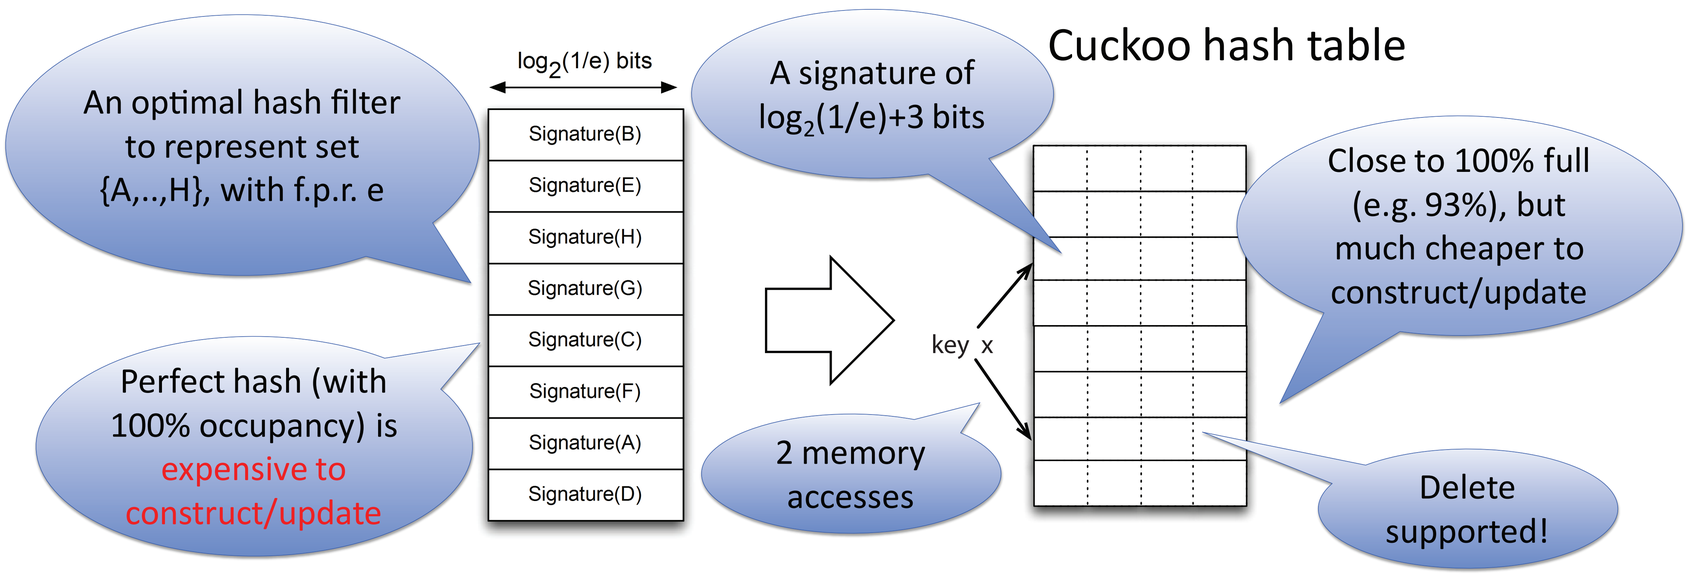

HASH FILTER VS. BLOOM FILTER: WHEN TO USE

FALSE POSITIVE RATE V.S. BITS/KEY

- **Hash filter is more space efficient when f.p.r. < 2%**

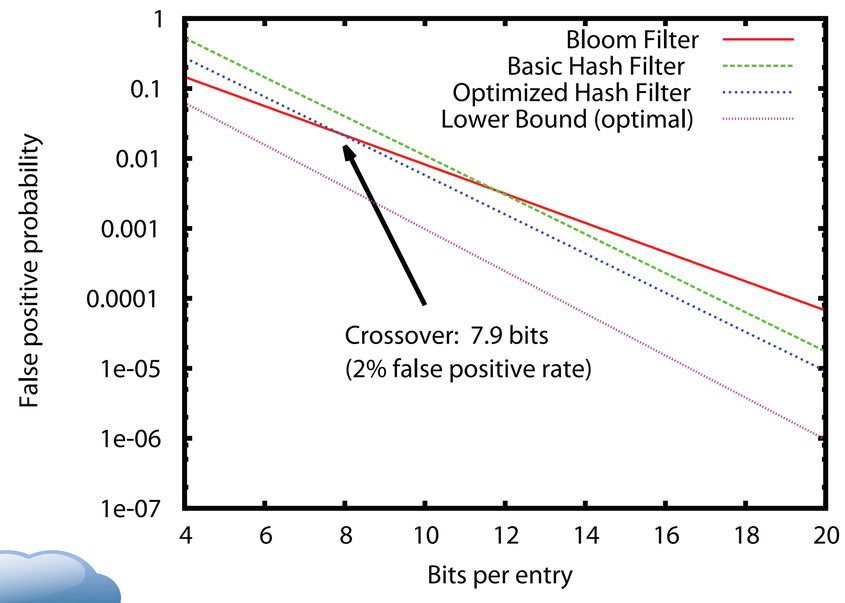

MEM ACCESSES V.S. BITS/KEY

- **When f.p.r. < 1%, hash filter wins in both metrics!**

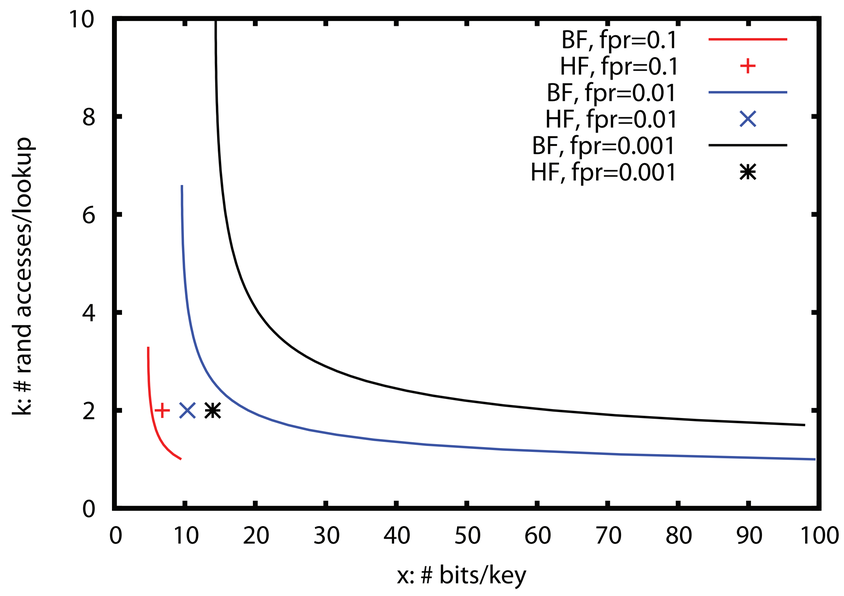

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · PRINCETON UNIVERSITY · UC Berkeley