

Analysis

Search takes $O(\log(N))$ time

Proof

- (intuition): *at each step, distance between query and peer-with-file reduces by a factor of at least 2*
- (intuition): after $\log(N)$ forwardings, distance to key is at most $2^m / 2^{m \cdot 2} = 2^m / N$
- Number of node identifiers in a range of $2^m / N$ is $O(\log(N))$ with high probability (why? SHA-1! and “Balls and Bins”)
So using *successors* in that range will be ok, using another $O(\log(N))$ hops

