

# Effects of skip length

- For small segment lengths (about 1000 bytes long), skipping data helps very little, even at very small skip lengths.
- For moderate lengths ( $10^4$  to  $10^5$  bytes long), very short skip distances give the largest expected performance increase.
- For long lengths (over  $10^6$  bytes long), short skip distances help, but longer skip distances are even better.

# Effects of skip length




Figure 2: The expected costs, relative to no skipping, of processing segments of varying lengths with varying skip sizes, based on TREC 2005 Efficiency Track queries.

# Multicore speedup?

- A/M (in memory) had 3.14x speedup on 4 cores. They had 2.35x speedup.
- Note: They were 13x faster than A/M (not in memory) when they started. They are now only about 9.7x faster on 4 cores, assuming similar speedups for both implementations of A/M.