



Ανάκληση Πληροφορίας

Information Retrieval

Διδάσκων –
Δημήτριος Κατσαρός

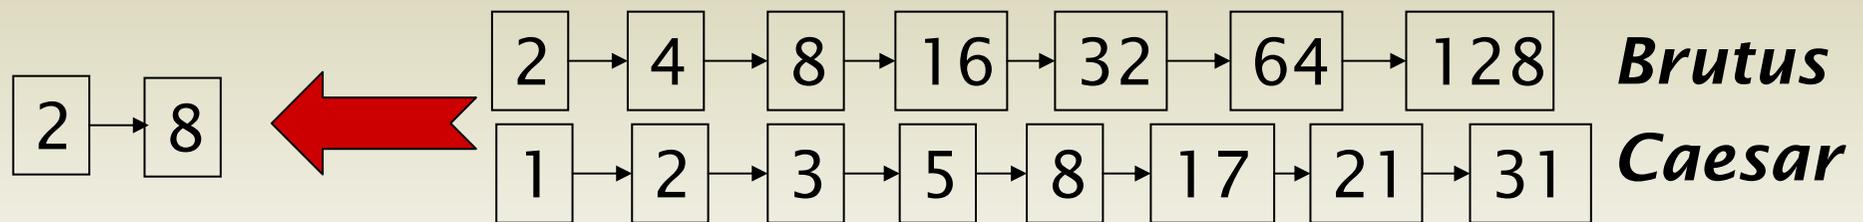


Faster postings merges:
Skip pointers



Recall basic merge

- Walk through the two postings simultaneously, in time linear in the total number of postings entries



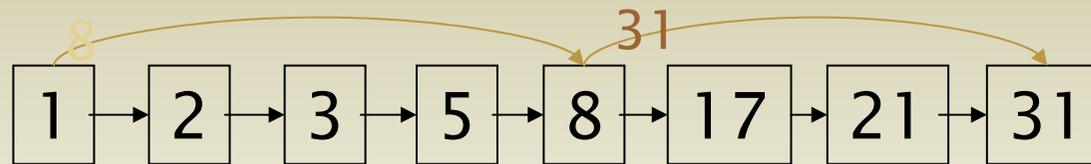
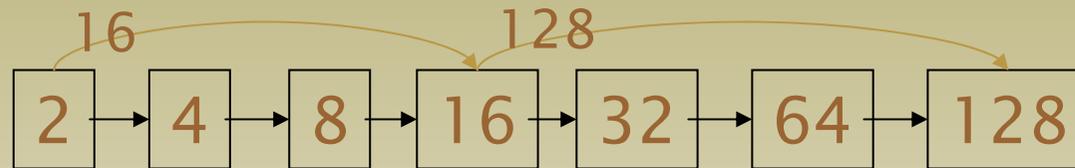
If the list lengths are m and n , the merge takes $O(m+n)$ operations.

Can we do better?

Yes, if index isn't changing too fast.



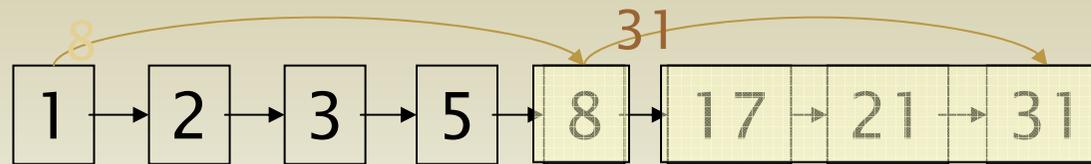
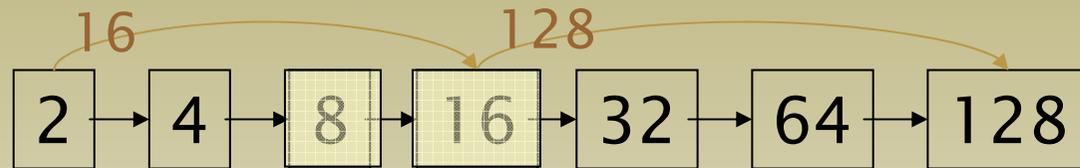
Augment postings with skip pointers (at indexing time)



- Why?
- To skip postings that will not figure in the search results.
- How?
- Where do we place skip pointers?



Query processing with skip pointers



Suppose we've stepped through the lists until we process **8** on each list.

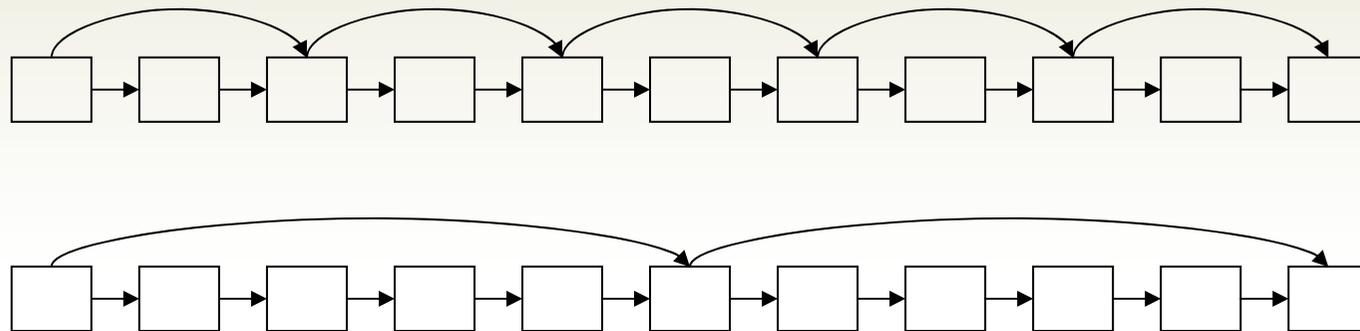
When we get to **16** on the top list, we see that its successor is **32**.

But the skip successor of **8** on the lower list is **31**, so we can skip ahead past the intervening postings.



Where do we place skips?

- Tradeoff:
 - More skips \rightarrow shorter skip spans \Rightarrow more likely to skip. But lots of comparisons to skip pointers.
 - Fewer skips \rightarrow few pointer comparison, but then long skip spans \Rightarrow few successful skips.





Placing skips

- Simple heuristic: for postings of length L , use \sqrt{L} evenly-spaced skip pointers.
- This ignores the distribution of query terms.
- Easy if the index is relatively static; harder if L keeps changing because of updates.
- This definitely used to help; with modern hardware it may not (Bahle et al. 2002)
 - The cost of loading a bigger postings list outweighs the gain from quicker in memory merging