

Lower Bounds for External Memory Dictionaries

Gerth Stølting Brodal Rolf Fagerberg

 BRICS

University of Aarhus

Aalborg University, February 12, 2003

Work presented at SODA, Jan 2003

Dictionary



- **Queries**

- membership
- predecessor / successor
- range queries ...

- **Updates**

- insertions
- deletions

Dictionary



- **Queries**

- membership
- predecessor / successor
- range queries ...


- **Updates**

- insertions
- deletions

This talk : Comparison based, membership, insertions


Dictionaries – Comparison Based

	Insert	Search
Balanced search trees	$O(\log N)$	$O(\log N)$




Dictionaries – Comparison Based

	Insert	Search
Balanced search trees	$O(\log N)$	$O(\log N)$
Adversary	∞	$\Omega(\log N)$

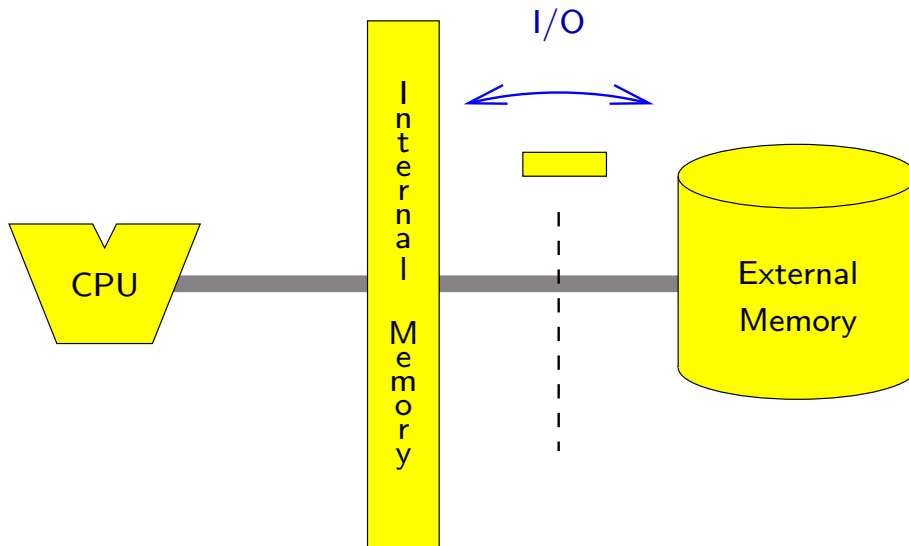


Dictionaries – Comparison Based

	Insert		Search
Balanced search trees	$O(\log N)$		$O(\log N)$
Adversary	∞		$\Omega(\log N)$
Borodin et al. 1982	$O(t)$	\Rightarrow	$N/2^{O(t)}$



External Memory Model



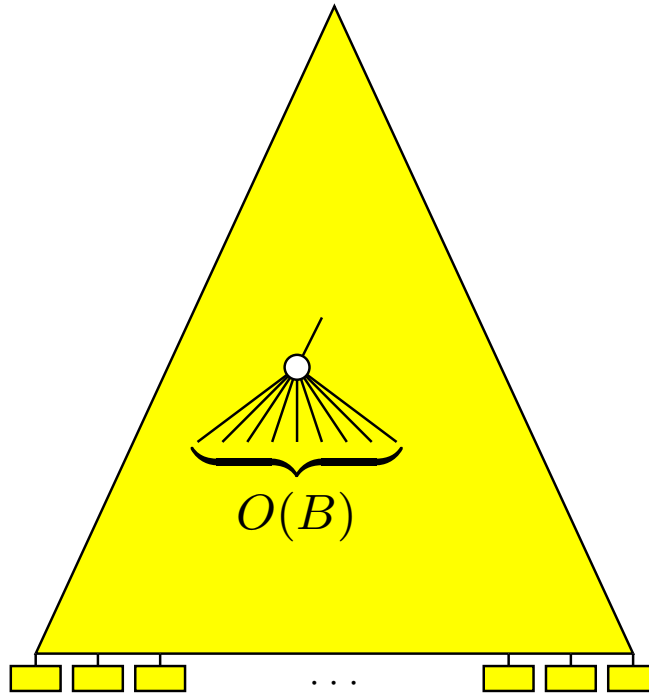
N = problem size
 M = memory size
 B = I/O block size

- One I/O moves B consecutive records from/to disk
- **Cost** : number of I/Os
- Elements can be copied and compared in internal memory

B-trees

– An External Memory Dictionary

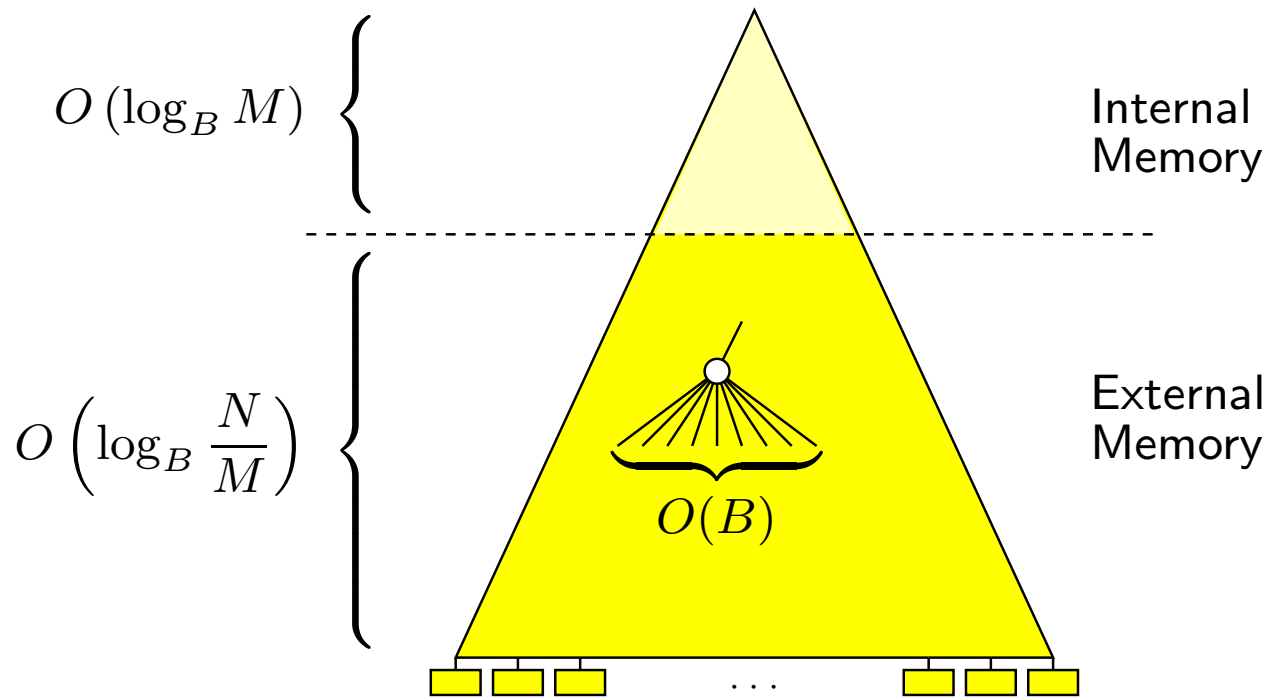
Bayer and McCreight 1972



B-trees

– An External Memory Dictionary

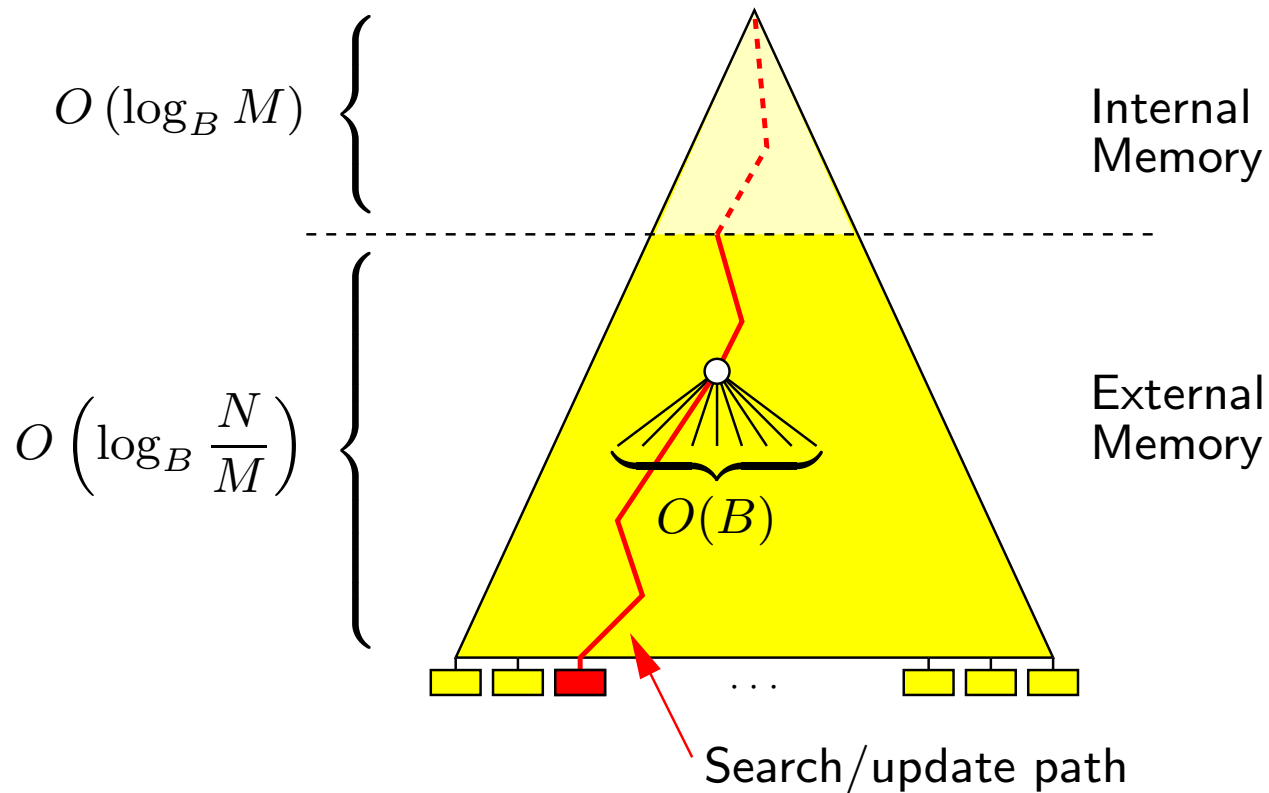
Bayer and McCreight 1972



B-trees

– An External Memory Dictionary

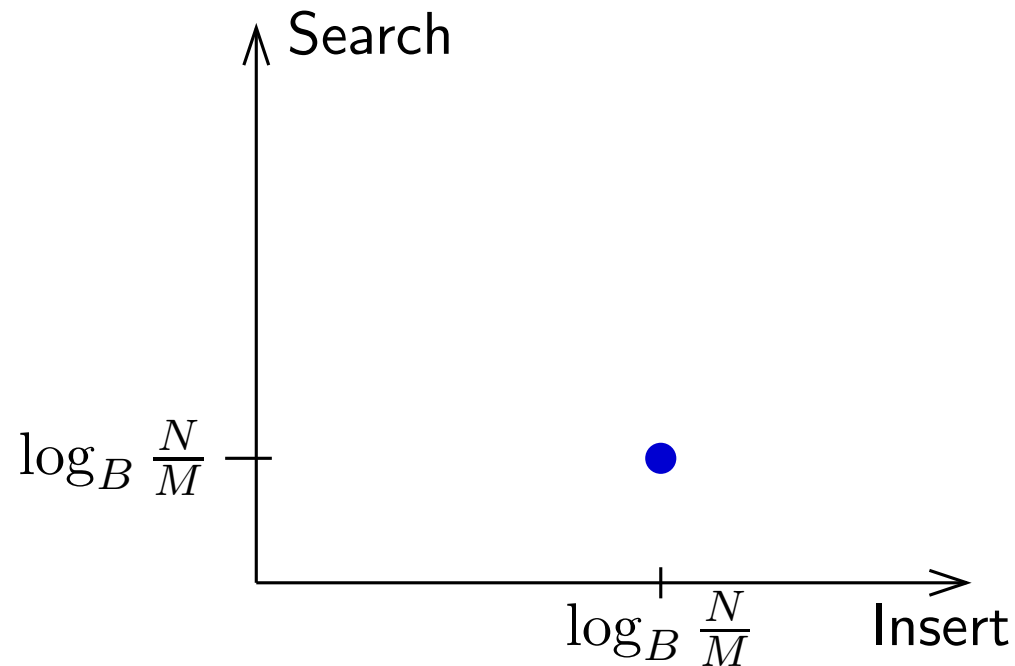
Bayer and McCreight 1972



$$\left. \begin{array}{l} \text{Insert} \\ \text{Search} \end{array} \right\} O\left(\log_B \frac{N}{M}\right) \text{ I/Os}$$


Dictionaries – External Memory

	Insert	Search
B-trees	$O(\log_B \frac{N}{M})$	$O(\log_B \frac{N}{M})$




Dictionaries – External Memory

	Insert	Search
B-trees	$O(\log_B \frac{N}{M})$	$O(\log_B \frac{N}{M})$
Adversary	∞	$\Omega(\log_B \frac{N}{M})$




Dictionaries – External Memory

	Insert	Search
B-trees	$O(\log_B \frac{N}{M})$	$O(\log_B \frac{N}{M})$
Adversary	∞	$\Omega(\log_B \frac{N}{M})$





Dictionaries – External Memory

	Insert	Search
B-trees	$O(\log_B \frac{N}{M})$	$O(\log_B \frac{N}{M})$
Adversary	∞	$\Omega(\log_B \frac{N}{M})$





Question addressed in this talk

Comparisons vs. I/Os



Comparisons vs. I/Os



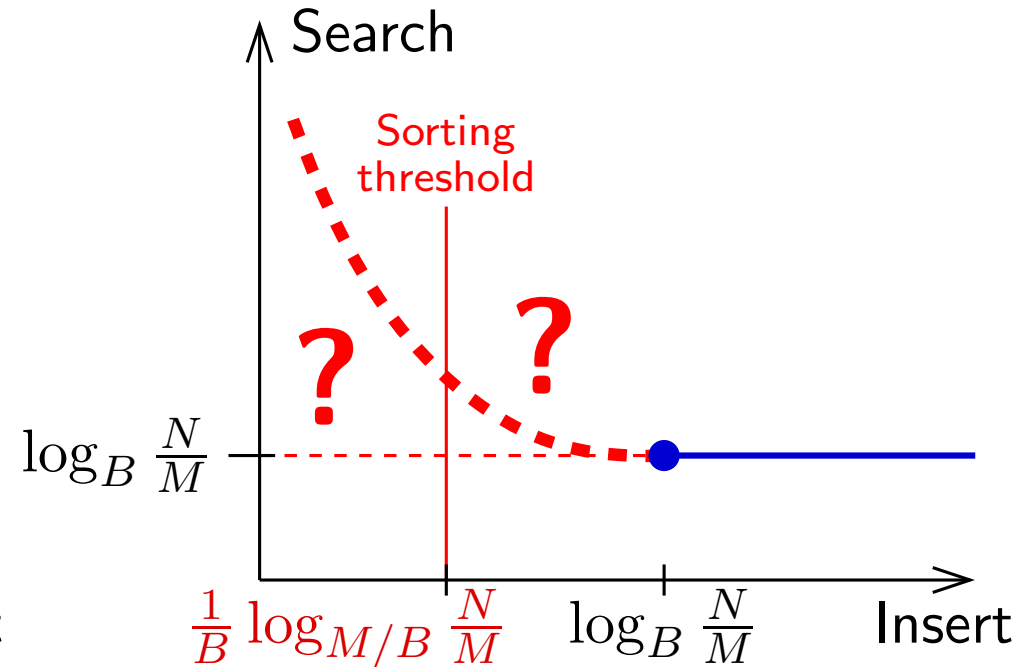
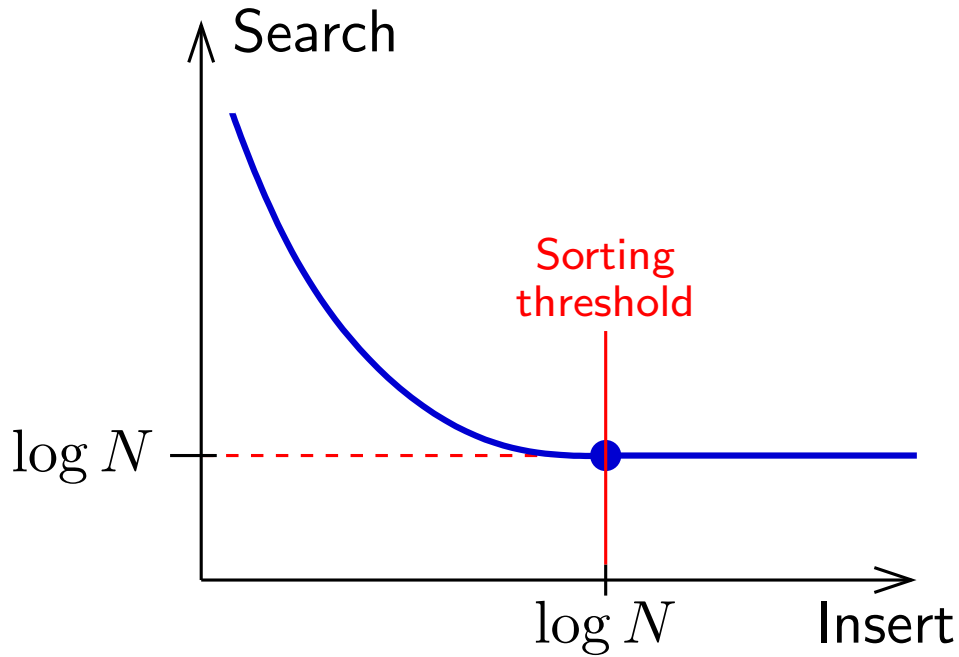
Sorting

Comparisons $\Theta(N \log N)$

I/Os $\Theta\left(\frac{N}{B} \log_{M/B} \frac{N}{M}\right)$

Aggarwal and Vitter 1988

Comparisons vs. I/Os




Sorting

Comparisons $\Theta(N \log N)$

I/Os $\Theta\left(\frac{N}{B} \log_{M/B} \frac{N}{M}\right)$

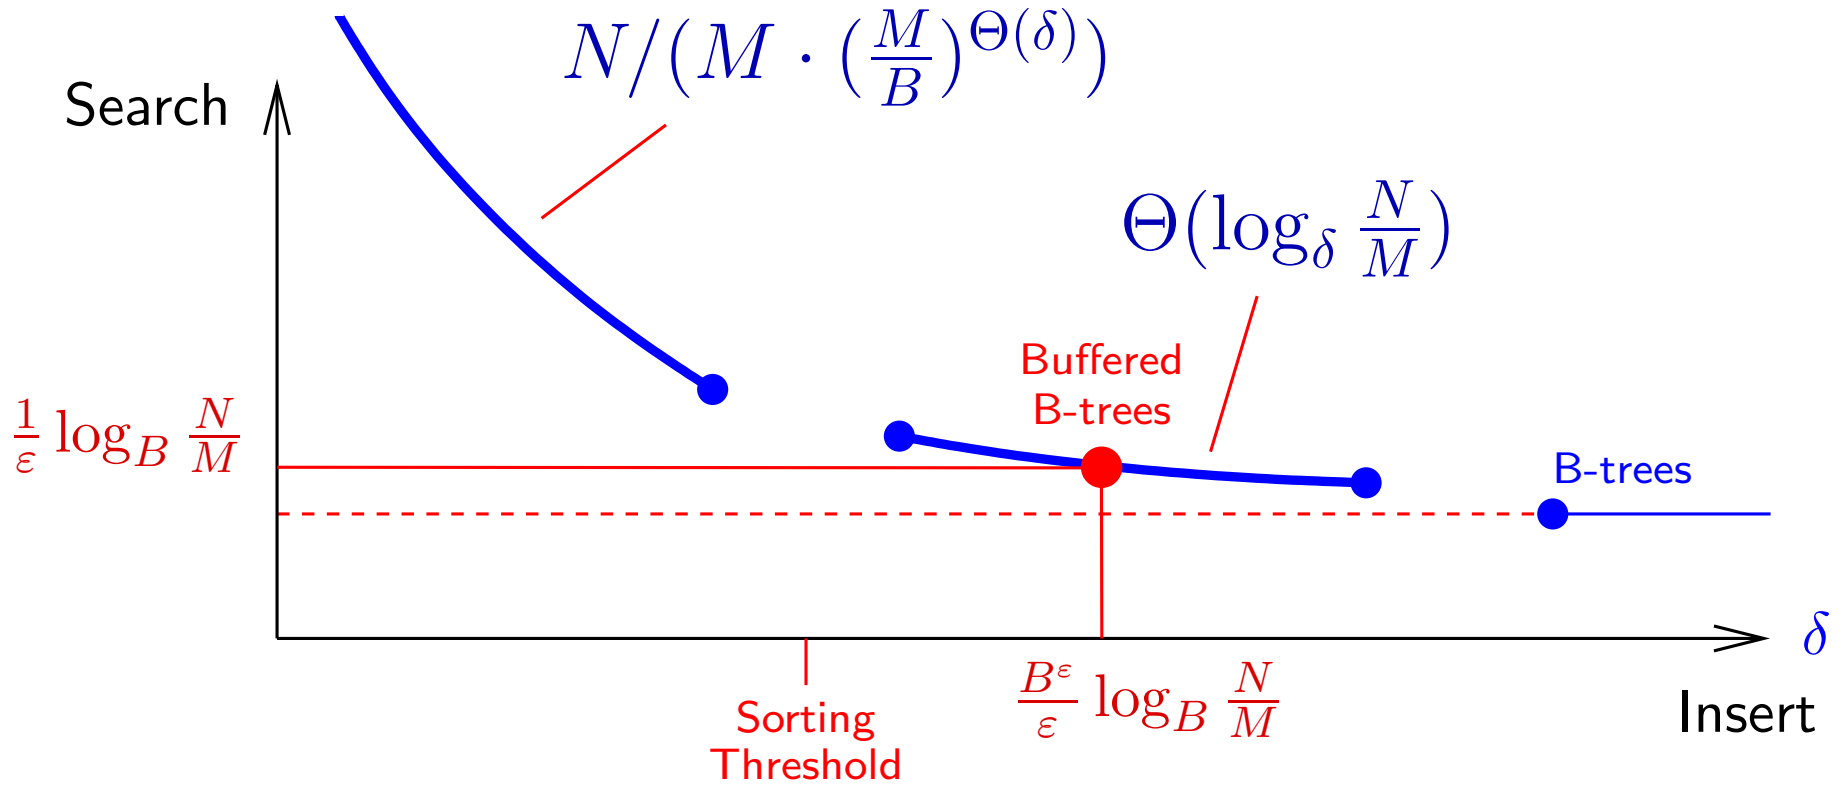
Aggarwal and Vitter 1988

Results




δ = number of I/Os for B insertions

Results



δ = number of I/Os for B insertions


Results



δ = number of I/Os for B insertions

Buffered B-trees

– how to speedup B-tree updates by a factor $\varepsilon B^{1-\varepsilon}$



Searches

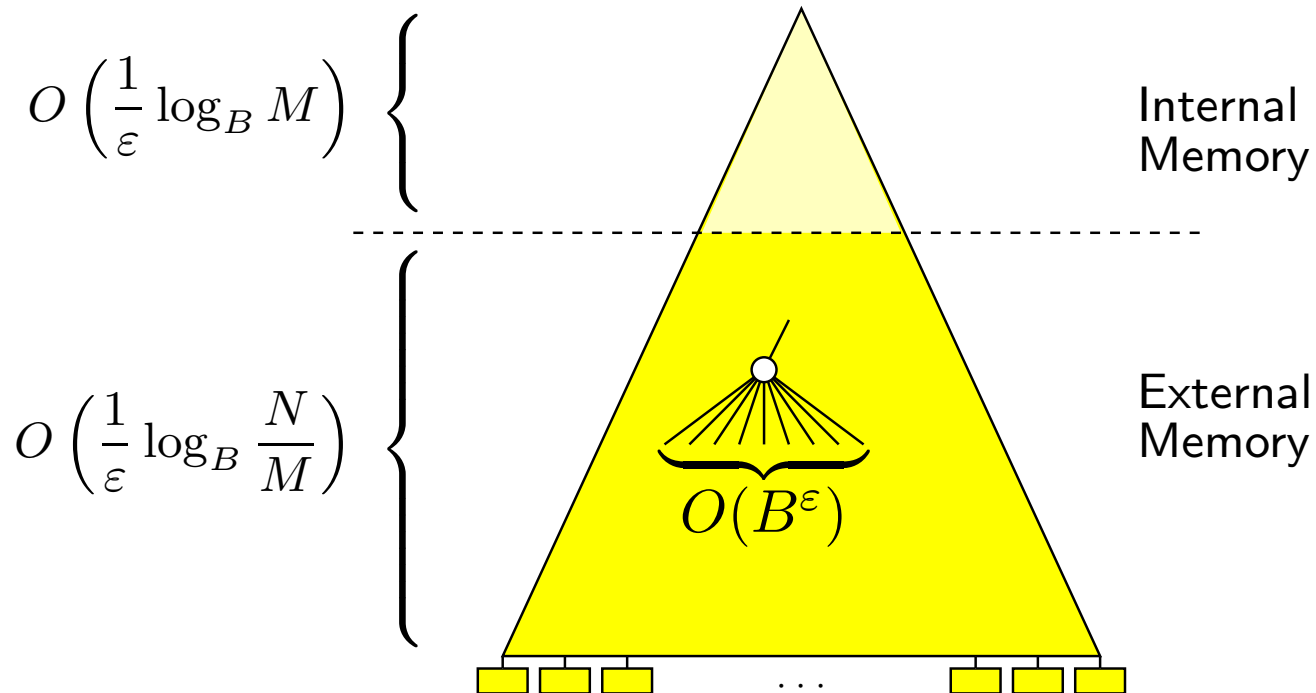
$$O\left(\frac{1}{\varepsilon} \log_B \frac{N}{M}\right)$$

B insertions

$$O\left(\frac{B^\varepsilon}{\varepsilon} \log_B \frac{N}{M}\right)$$

Buffered B-trees

– how to speedup B-tree updates by a factor $\varepsilon B^{1-\varepsilon}$



- B-tree with degree $\Theta(B^\varepsilon)$

Searches


$$O\left(\frac{1}{\varepsilon} \log_B \frac{N}{M}\right)$$

B insertions

$$O\left(\frac{B^\varepsilon}{\varepsilon} \log_B \frac{N}{M}\right)$$

Buffered B-trees

– how to speedup B-tree updates by a factor $\varepsilon B^{1-\varepsilon}$



- B-tree with degree $\Theta(B^\varepsilon)$
- Buffers of $O(B)$ delayed insertions

Searches


$$O\left(\frac{1}{\varepsilon} \log_B \frac{N}{M}\right)$$

B insertions

$$O\left(\frac{B^\varepsilon}{\varepsilon} \log_B \frac{N}{M}\right)$$

Buffered B-trees

– how to speedup B-tree updates by a factor $\varepsilon B^{1-\varepsilon}$



- B-tree with degree $\Theta(B^\varepsilon)$
- Buffers of $O(B)$ delayed insertions

Searches

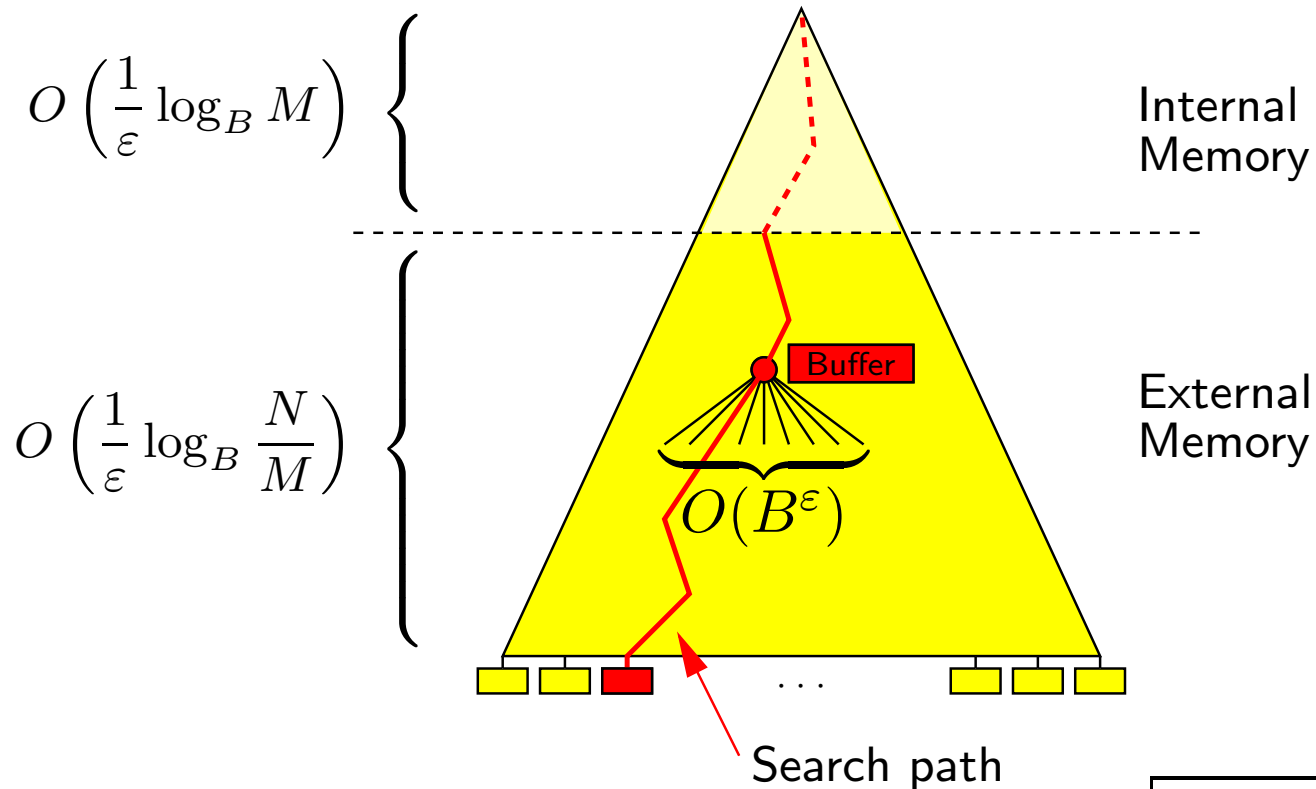
$$O\left(\frac{1}{\varepsilon} \log_B \frac{N}{M}\right)$$

B insertions

$$O\left(\frac{B^\varepsilon}{\varepsilon} \log_B \frac{N}{M}\right)$$

Buffered B-trees

– how to speedup B-tree updates by a factor $\varepsilon B^{1-\varepsilon}$



- B-tree with degree $\Theta(B^\varepsilon)$
- **Buffers** of $O(B)$ delayed insertions
- On buffer overflow move $O(B^{1-\varepsilon})$ elements to a child with one I/O

Searches


$$O\left(\frac{1}{\varepsilon} \log_B \frac{N}{M}\right)$$

B insertions

$$O\left(\frac{B^\varepsilon}{\varepsilon} \log_B \frac{N}{M}\right)$$


Lower Bound

– optimality of buffered B-trees



- Adversary online constructs S_1, \dots, S_K
- Goal: an i such that x_{ij} has not been in internal memory since end of S_j , for all $j = 1, \dots, K$
- x_{i1}, \dots, x_{iK} form an antichain, i.e. search requires $\geq K$ I/Os

Lower Bound (Cont.)



- Let \hat{I} be the indexes i where
 - $x_{ij} \in S_j$ but is not in internal memory at end of S_j
 - $x_{i1}, \dots, x_{i(j-1)}$ have not been read into internal memory by the $\delta \frac{|S_j|}{B}$ I/Os during the insertion of S_j
- Construct $I \subset \hat{I}$ such that all blocks in external memory contain $O(\frac{B}{\delta})$ elements x_{ij} where $i \in I$
 - Existence by randomized sampling with probability $O(1/\delta)$ and Chernoff bounds, provided $B/\delta = \Omega(\log N)$
- Let $x_{i(j+1)} \in S_{j+1}$ iff $i \in I$

$$K = \Theta(\log_{\delta} \frac{N}{M})$$

Lower Bound — Below Sorting

Insert		Search
$\frac{\delta}{B}$	\Rightarrow	$\frac{N}{M \cdot \left(\frac{M}{B}\right)^{\Theta(\delta)}}$

Lower Bound — Below Sorting

Insert		Search
$\frac{\delta}{B}$	\Rightarrow	$\frac{N}{M \cdot \left(\frac{M}{B}\right)^{\Theta(\delta)}}$

- W.l.o.g. memory and each block totally ordered after each I/O

$$N \log M + \delta \frac{N}{B} \left(B \log \frac{M}{B} \right) \text{ comparisons}$$

Insert in internal memory

Merging a block with internal memory

Lower Bound — Below Sorting

Insert		Search
$\frac{\delta}{B}$	\Rightarrow	$\frac{N}{M \cdot \left(\frac{M}{B}\right)^{\Theta(\delta)}}$

- W.l.o.g. memory and each block totally ordered after each I/O

$$N \log M + \delta \frac{N}{B} \left(B \log \frac{M}{B} \right) \text{ comparisons}$$

↑
Insert in internal memory


↑
Merging a block with internal memory

- Antichain of size (Borodin et al. 1982 / Dillworth's lemma)

$$\frac{N}{2^{\log M + \delta \log \frac{M}{B}}} = \frac{N}{M \cdot \left(\frac{M}{B}\right)^\delta}$$

of which all elements except one are in distinct blocks

Conclusion



Conclusion

