

Cache-Oblivious Dynamic Dictionaries with Update/Query Tradeoff

Gerth Stølting Brodal

Erik D. Demaine

Jeremy T. Fineman

John Iacono

Stefan Langerman

J. Ian Munro

Result presented at SODA 2010

Dynamic Dictionary


Search(k)

Insert(e)

Delete(k)

I/O Model


[Aggarwal, Vitter 88]



Cost: the number of *block transfers* (I/Os)

Cache-Oblivious Algorithms

[Frigo, Leiserson, Prokop, Ramachandran 99]



- Algorithms not parameterized by M or B
- Analyze in *ideal-cache model* – I/O model, except optimal replacement policy is assumed

Cache-Oblivious Dynamic Dictionaries


Cache-Aware	Search	Insert
B-tree [BM72]	$\alpha(\log_B M)$	$\alpha(\log_B M)$
Buffered B-tree [BF03]	$\alpha(1/\varepsilon)\log_B M$	$\alpha(1/\varepsilon B^{1-\varepsilon})\log_B M^*$

Cache-Oblivious	Search	Insert
CO B-tree [BDF-00, BDIW04, BFJ02]	$\alpha(\log_B M)$	$\alpha(\log_B N + \dots)$
COLA [BFF-CFKN07]	$\alpha(\log_2 M)$	$\alpha(1/B)\log_2 M^*$
Shuttle Tree [BFF-CFKN07]	$\alpha(\log_B M)$	$\alpha(1/B^{\Omega(1/(\log \log B)^2)})\log_B N + \dots)^*$
xDict [this paper]	$\alpha(1/\varepsilon)\log_B M$	$\alpha(1/\varepsilon B^{1-\varepsilon})\log_B M^{*\dagger}$

* amortized

† assumes $M = \Omega(B^2)$

Building an xDict ($\epsilon = 1/2$)



$\lg \lg N$ x-boxes of squaring capacities

Insert: insert into smallest box

- When a box reaches capacity, **Flush** it and **Batch-Insert** into the next box
- $\alpha(1/\sqrt{B}) \log_B x$ cost is dominated by largest box
→ $\alpha(1/\sqrt{B}) \log_B N$

Search: search in each x-box

- $\alpha(\log_B x)$ cost is dominated by largest box $\alpha(\log_B N)$


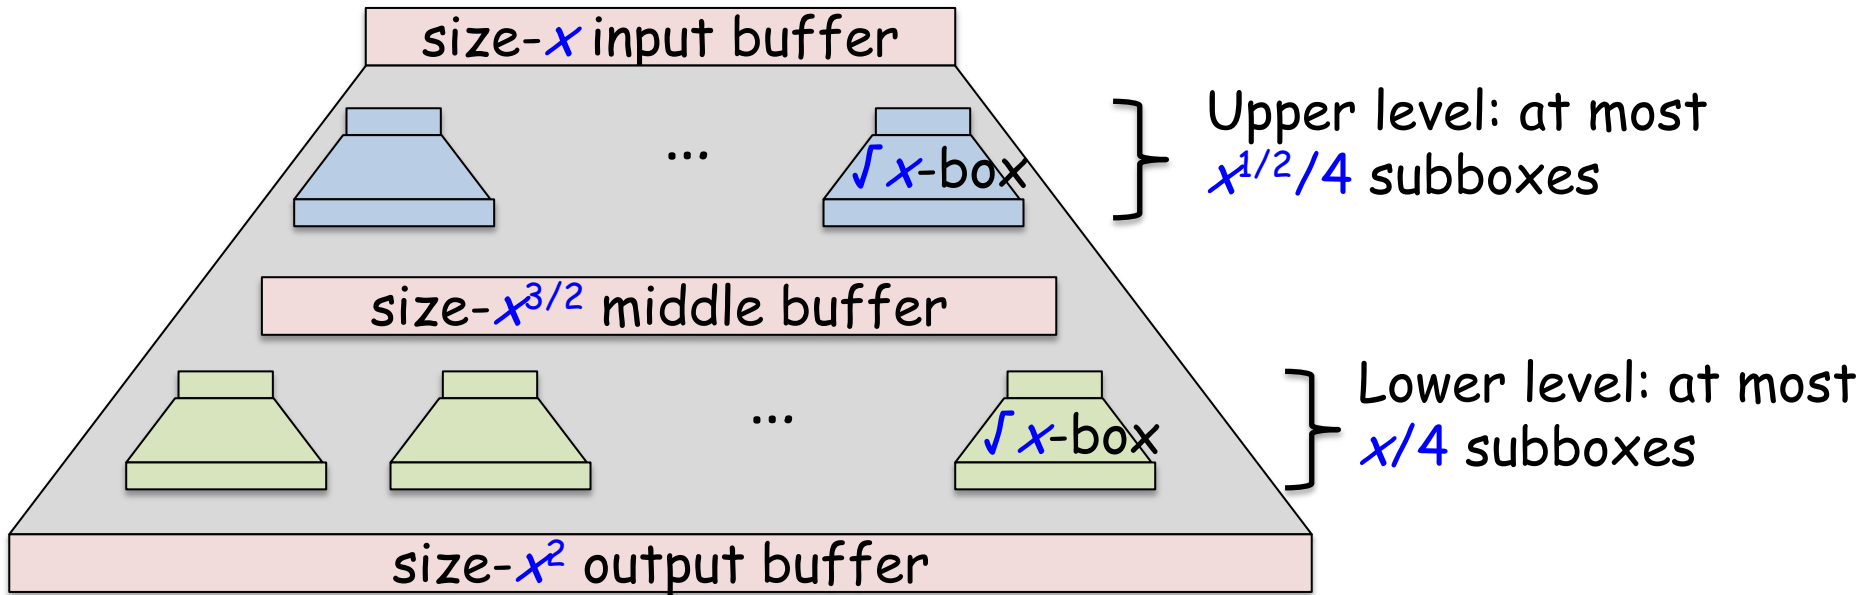
x-Box = dictionary with capacity x^2

Batch-Insert(D, A): insert $\Theta(x)$ presorted objects
– cost $\alpha(1/\sqrt{B})\log_B x$ per element

Search(D, k):
– cost is $\alpha(\log_B x)$

Flush(D): produce a size- x^2 sorted array A containing all the elements in the x -box D
– cost is $\alpha(1/B)$ per element

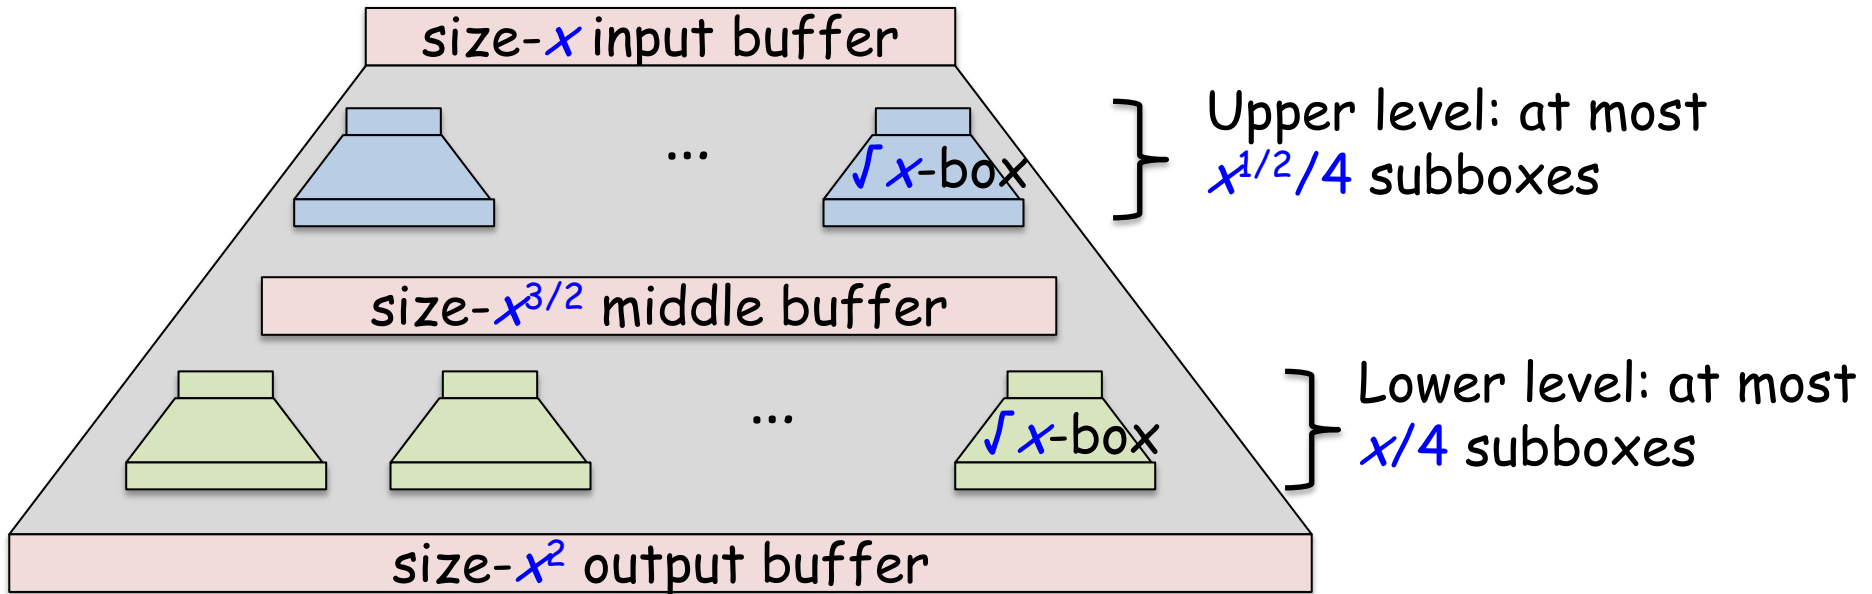
Recursive x-Box



subboxes stored contiguously
in arbitrary order

Unused (currently empty)
subboxes are preallocated

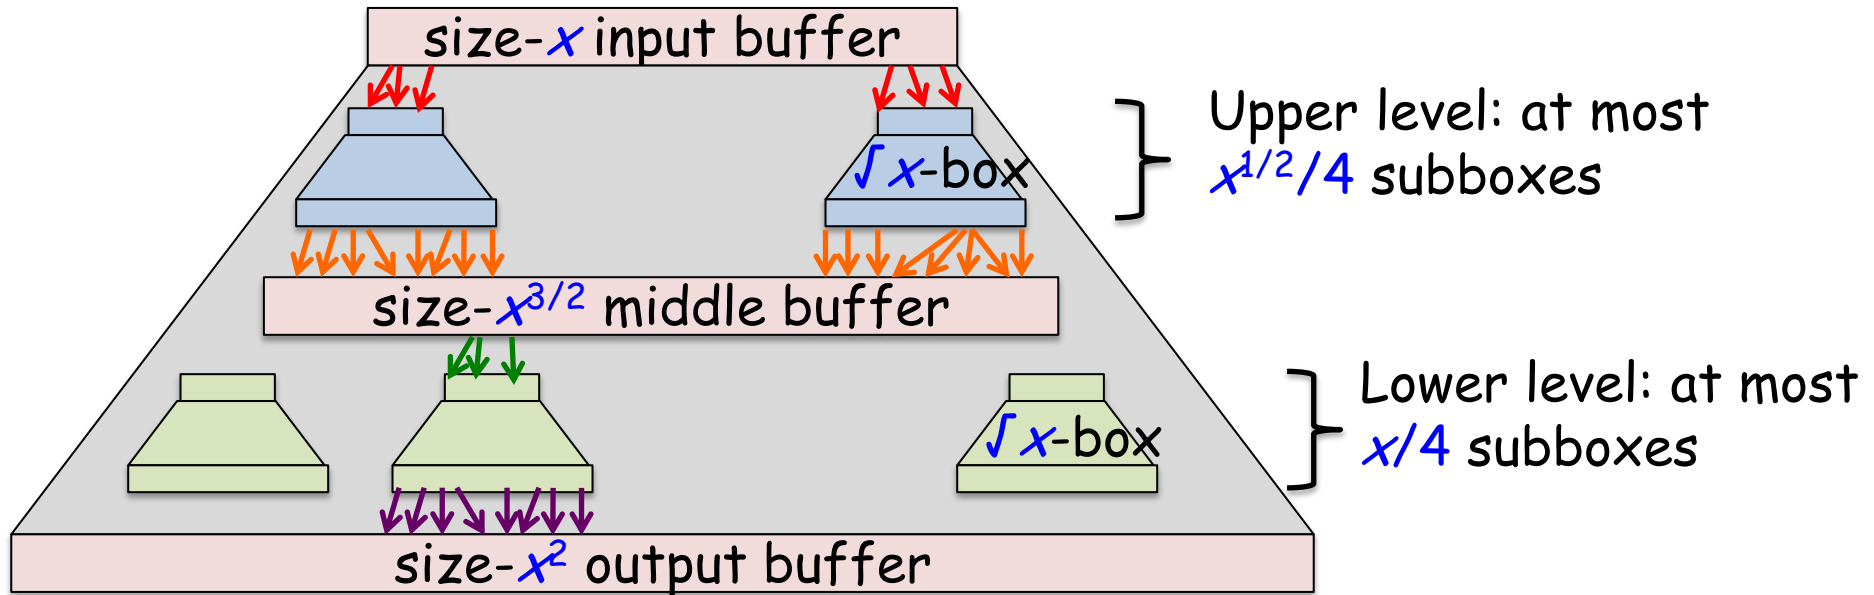
x-Box Space Usage



Theorem: An x -Box uses at most cx^2 space

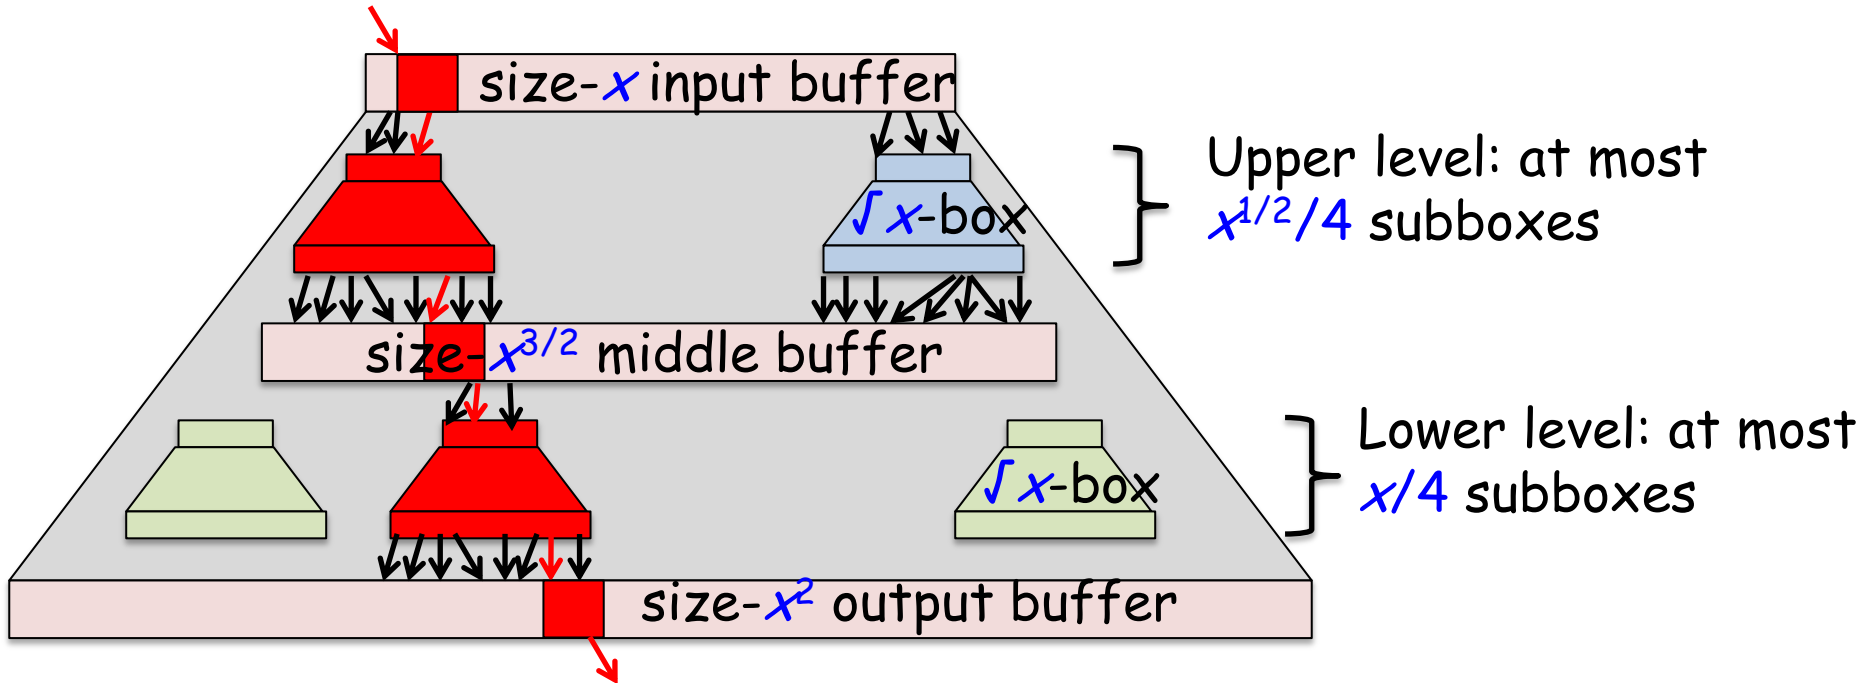
(within constant factor of capacity/output buffer)

Fractional Cascading within x -Box



Propagate samples upwards + Lookahead pointers

Searching in an x -Box



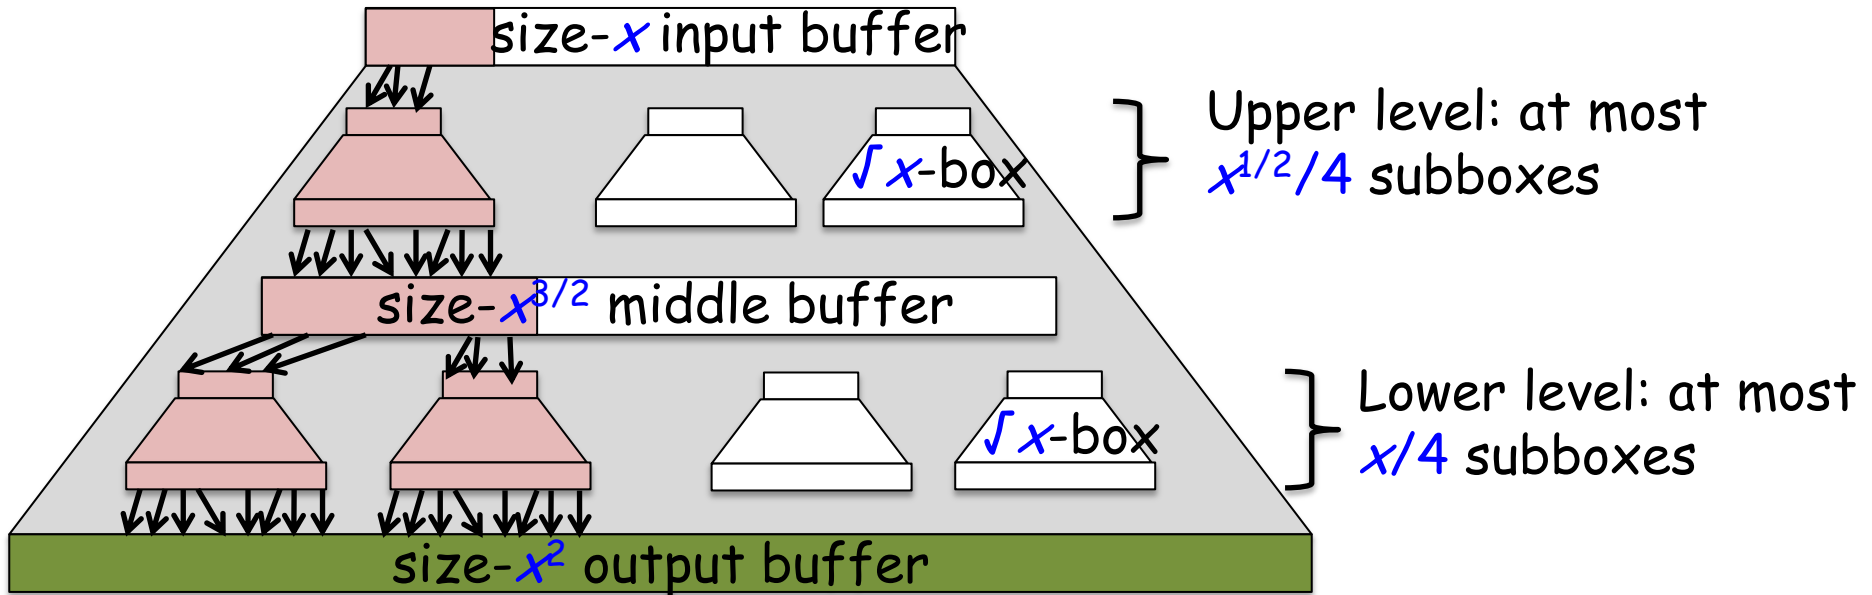
Describe searches by the recurrence

$$S(x) = 2S(\sqrt{x}) + O(1)$$

with base case $S(<\sqrt{B}) = 0$


Solves to $O(\log_B N)$

Flush




- Moves all real elements to the output buffer in sorted order.
- Lookahead pointers are rebuilt to facilitate searches. Most subboxes remain empty.

Batch-Insert




1. Merge sorted input into input buffer.

Batch-Insert




1. Merge sorted input into input buffer.
2. If input buffer is "full enough," Batch-Insert into upper-level subboxes (in chunks of $\Theta(\sqrt{x})$)


Batch-Insert




1. Merge sorted input into input buffer.
2. If input buffer is "full enough," Batch-Insert into upper-level subboxes (in chunks of $\Theta(\sqrt{x})$)
3. Whenever a subbox is near capacity, Flush it, then split it into two subboxes



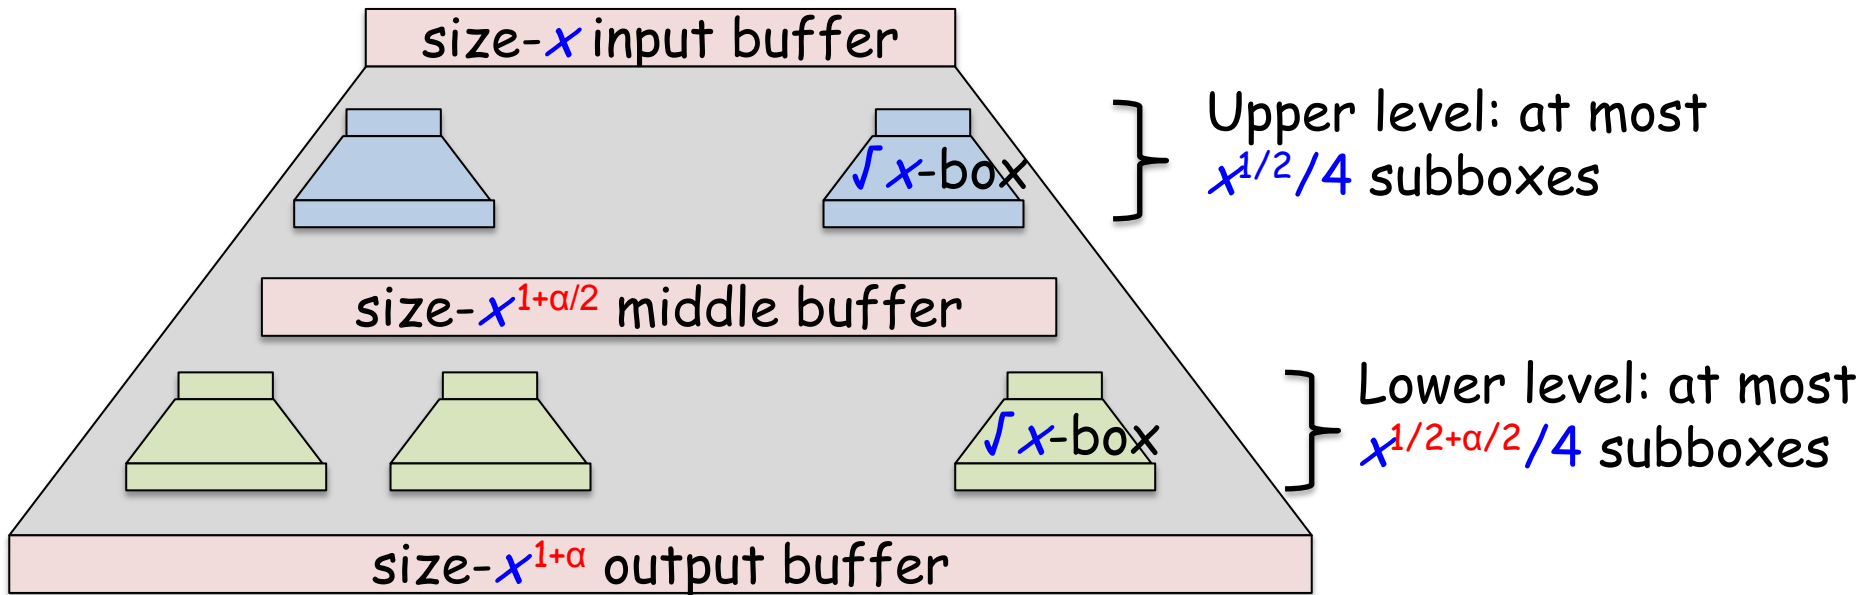
Batch-Insert



1. Merge sorted input into input buffer.
2. If input buffer is "full enough," Batch-Insert into upper-level subboxes (in chunks of $\Theta(\sqrt{x})$)
3. Whenever a subbox is near capacity, Flush it, then split it into two subboxes
4. If no empty subboxes remain, Flush all of them and merge output buffers into middle buffer.



Generalizing to $O((1/\varepsilon B^{1-\varepsilon}) \log_B N)$



Parameterize by $0 < \alpha \leq 1$, where $\alpha = \varepsilon/(1-\varepsilon)$



$1/\varepsilon$ overhead comes from geometric sum in xDict

Results Summary

Cache-Aware	Search	Insert
B-tree [BM72]	$\alpha(\log_B M)$	$\alpha(\log_B M)$
Buffered B-tree [BF03]	$\alpha(1/\varepsilon)\log_B M$	$\alpha(1/\varepsilon B^{1-\varepsilon})\log_B M^*$

Cache-Oblivious	Search	Insert
CO B-tree [BDF-00, BDIW04, BFJ02]	$\alpha(\log_B M)$	$\alpha(\log_B N + \dots)$
COLA [BFF-CFKN07]	$\alpha(\log_2 M)$	$\alpha(1/B)\log_2 M^*$
Shuttle Tree [BFF-CFKN07]	$\alpha(\log_B M)$	$\alpha(1/B^{\Omega(1/(\log \log B)^2)})\log_B N + \dots)^*$
xDict [this paper]	$\alpha(1/\varepsilon)\log_B M$	$\alpha(1/\varepsilon B^{1-\varepsilon})\log_B M^{*\dagger}$

* amortized

† assumes $M = \Omega(B^2)$