

Parallel computing - Parallel Algorithms João Luís Sobral (jls@di ...) 25/Nov/2021





### **Parallel Algorithms**

- Traditional algorithm analysis: Big O notation
  - Analysis of the number of operations per datum (n)
    - Matriz multiplication:  $O(n^3) O(n^2)$  (n=number of elements of C)
    - Sorting (n keys)
      - Bad sorting (burte force): O(n<sup>2</sup>)
      - Better algorithm: n log<sub>2</sub>(n)
      - Best: ([k]n)
    - Insertion on a data structure (n insertions):
      - Sorted linked list O(n<sup>2</sup>)
      - Binary tree: O(n log<sub>2</sub>(n))
      - Hash table O(n)





### **Parallel Algorithms**

• Arithmetic Intensity (operations per byte loaded)





#### • [Sequential] sorting algorithms

Method	Complexity (average)	Description
Insertion sort	n <sup>2</sup>	Insert elements into a sorted list
Bubble sort	n <sup>2</sup>	Compare (and swap) successive elements
Quicksort	n log n	Recursively sort elements less/greater than a given pivot
Merge sort	n log n	Successively merge sorted sub-lists starting from lists with one element
Heap sort	n log n	Insert elements into a binary heap
Radix sort	n d	Sort elements digit by digit (d) (two variants: MSD and LSD)







- Locality: Radix MSD vs Radix LSD
  - Both are based on key-indexed sorting
    - Example for 3-digit Strings (figure: I<sup>st</sup> digit step on MSD)
      - 1. Count frequencies of each letter using key as index
      - 2. Compute frequency cumulates
      - 3. Access cumulates using key as index to find record positions.
      - 4. Copy back into original array

```
int N = a.length;

int[] count = new int[R];

for (int i = 0; i < N; i++)

count[a[i]+1]++;

compute

cumulates \rightarrow for (int k = 1; k < 256; k++)

count[k] += count[k-1];

move

records \rightarrow for (int i = 0; i < N; i++)

temp[count[a[i]++]] = a[i]

copy back \rightarrow for (int i = 0; i < N; i++)

a[i] = temp[i];
```



JLSobral, Parallel Computing, MEI, UMinho, 2021/22



- Locality: Radix LSD (e.g. string with 3 digits)
  - D steps through the full data: DxN data moves





- Locality: Radix MSD(e.g. string with 3 digits)
  - Partition data in K-sets- I step through the full data: IxN (global) + local data moves





- Locality: quick-sort vs heap-sort
  - Regular VS irregular data references



Both perform *n x log n* steps through the data but heapsort shows more irregular accesses to data indexes

JLSobral, Parallel Computing, MEI, UMinho, 2021/22







Locality of reference in sorting algorithms

Method	Locality of reference	Improvements
Quicksort	Good spatial locality + bad temporal locality on initial stages	Initial set partioning using k keys
Mergesort	Good spatial locality + bad temporal locality on final merge stages	Single merge when data exceeds cache size
Heap sort	Bad	Cache aware trees + d- fan-out
Insertion sort	Bad	
Radix sort	Good when MSD first (only when processing LSDs)	Reduce the number passes through data



Parallelism in sorting algorithms

Method		
Quicksort	Sort sub-lists in parallel / start with p lists	
Merge-sort	Merge p lists in parallel	
Heap sort	???	
Insertion sort	???	
Radix sort	Sort set of digits in parallel	

#### Parallel Sorting (on distributed memory)

- Design issues:
  - Keys are initially distributed over processors
    - Intermediate stage for other parallel algorithms
  - Data properties
    - Partially-sorted data? (not in the scope of this lecture)
  - Exploitable parallelism
    - Merger-based
    - Splitter-based
  - Efficiency considerations for parallelism
    - Data movements [across processors]
    - Load balancing
    - Avoid idle time (e.g. sequential phases)



- Parallel Merge-Sort
  - Locally sort each set
  - Exchange sets among processors



- Only effective when  $n/p \sim I$
- Extensive data movements: when n/p>>I



- Parallel quicksort (simplified)
  - Master selects and broadcasts pivot key
  - Each process locally splits using the pivot
    - Each processes holds smaller and greater partitions
  - Divide processors into smaller and greater sets
    - Send data to one processor on the other set
  - Repeat the processes until #sets = #p
  - Locally sort on each process p
  - Complexity:
    - Requires log p communication steps

JLSobral, Parallel Computing, MEI, UMinho, 2021/22





- Parallel radix sort
  - Each processor is responsible by a subset of digit values
  - Sort and count the number of digit values
  - All-reduce the total number of digits
  - Send keys to the processor responsible for each digit range
  - Repeat for the next digit
  - Complexity:
    - LSD D communication steps
    - MSD I communication step



- Parallelism in sorting algorithms
  - Sampling based
    - Split data into P sets using p-1 splitters
    - Each processor acts upon a local set
    - Minimizes data movements
  - Sampling alternatives
    - Regular sampling (p\*(p-1) keys)
      - Not effective for large p
    - Random sampling
    - Histogram sampling

#### Parallel Sorting by Regular Sampling

Divide the set into *p* disjoint sets and locally order each set

1.

4.

0

0

- Applies a local QuickSort Selects *p-1* local samples that uniformly divide each set into p subsets
- Order p\*(p-1) samples and select best p-1 pivot keys
- Partition each set using the p-1 pivot keys
  - Merge *p\*p* sets
    - Processor *i* merges the *i* partition

