Stanford University

Radix Sort

1 Radix Sort

1.1 Some background

Before detailing the Radix Sort algorithm, we mention a couple of properties of this sort:

- It is not a comparison sort. This means that contrary to Merge Sort, Quick Sort, Insertion Sort, ..., the sorting does not only rely on a comparison operator (i.e., the less-or-equal operator), instead it compares each element bit-wise.

- It is a stable sort. It means that the sorting maintains the relative order of elements that have equal values.

- The sorting is not done in place. This implies that we will need a temporary array to store partial results.

- The complexity of the algorithm is $\mathcal{O}(kn)$ where k is the average element length.

You may find this online video useful. There are many online tutorials that you can check as well.

1.2 The algorithm

Radix Sort[1] sorts an array of elements in several passes. To do so, it examines, starting from the least significant bit, a group of `numBits` bits, sorts the elements according to this group of bits, and proceeds to the next group of bits.
More precisely:

1. Select the number of bits `numBits` you want to compare per pass.

2. Fills a histogram with numBuckets $= 2^{\text{numBits}}$ buckets, i.e., make a pass over the data and count the number of elements in each bucket.

3. Reorder the array to take into account the bucket to which an element belongs.

4. Process the next group of bits and repeat until you have dealt with all the bits of the elements (in our case 32 bits).

For instance let us say we want to sort:

$$\texttt{keys} = \boxed{0010 \mid 1011 \mid 0111 \mid 0000 \mid 0101 \mid 1111 \mid 1101 \mid 1001}$$

Step 1: We choose `numBits` = 2.
Step 2:

$$\texttt{buckets} = \underbrace{\boxed{1 \mid 3 \mid 1 \mid 3}}_{\substack{00 \quad 01 \quad 10 \quad 11}}$$

---

[1]In this programming assignment, we will implement LSD (least significan digit) Radix Sort.

Step 3:

$$\text{new\_keys} = \underbrace{\boxed{0000}\,\boxed{0101}}_{00}\,\underbrace{\boxed{1101}\,\boxed{1001}}_{01}\,\underbrace{\boxed{0010}\,\boxed{1011}}_{10}\,\underbrace{\boxed{0111}\,\boxed{1111}}_{11}$$

Step 4: Repeat with the two most significant bits.

1.3 A detailed example

We want to sort the following array:

$$\text{keys} = \boxed{001\,\,101\,\,011\,\,000\,\,010\,\,111\,\,110\,\,100}$$

Let us say that numBits $= 1$, i.e., we process one bit at a time.

First pass

The first thing we do is computing the histogram: in our case we will have numBuckets $= 2^{\texttt{numBits}}$
$= 2$ and:

$$\text{histogramRadixFrequency} = \boxed{4\,\,4}$$

We scan this histogram (i.e., we create a cumulative sum of the elements, starting at zero and ignoring the last element since it's equal to the number of elements):

$$\text{exScanHisto} = \boxed{0\,\,4}$$

The next step is to fill the temporary array. To do this, we need to keep track of the local offset in each of the bucket (the local offset will be used to compute the global offset, i.e., the position in the (partially) sorted array). We do this using:

$$\text{localOffsets} = \boxed{0\,\,0}$$

We now can fill the (partially) sorted array by reading keys and placing the elements in temp_keys (this step is called scattering). We start with:[2]

$$\text{temp\_keys} = \boxed{\,\,\,\,\,\,\,\,\,\,\,\,\,\,}$$

After reading the first element in keys we have:

$$\begin{aligned}\text{temp\_keys} &= \boxed{\,\,\,\,\,\,\,\,001\,\,\,\,\,\,}\\ \text{localOffsets} &= \boxed{0\,\,1}\end{aligned}$$

After the second:

$$\begin{aligned}\text{temp\_keys} &= \boxed{\,\,\,\,\,\,\,\,001\,\,101\,\,\,\,}\\ \text{localOffsets} &= \boxed{0\,\,2}\end{aligned}$$

After the third:

$$\begin{aligned}\text{temp\_keys} &= \boxed{\,\,\,\,\,\,\,\,001\,\,101\,\,011\,\,}\\ \text{localOffsets} &= \boxed{0\,\,3}\end{aligned}$$

---

[2]An empty cell means that the cell does not contain relevant information.

After the fourth:

$$\texttt{temp\_keys} \; = \; \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & & & & 001 & 101 & 011 & \\ \end{array}}$$

$$\texttt{localOffsets} \; = \; \boxed{\begin{array}{c|c} 1 & 3 \end{array}}$$

After the fifth:

$$\texttt{temp\_keys} \; = \; \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & 010 & & & 001 & 101 & 011 & \\ \end{array}}$$

$$\texttt{localOffsets} \; = \; \boxed{\begin{array}{c|c} 2 & 3 \end{array}}$$

After the sixth:

$$\texttt{temp\_keys} \; = \; \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & 010 & & & 001 & 101 & 011 & 111 \\ \end{array}}$$

$$\texttt{localOffsets} \; = \; \boxed{\begin{array}{c|c} 2 & 4 \end{array}}$$

After the seventh:

$$\texttt{temp\_keys} \; = \; \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & 010 & 110 & & 001 & 101 & 011 & 111 \\ \end{array}}$$

$$\texttt{localOffsets} \; = \; \boxed{\begin{array}{c|c} 3 & 4 \end{array}}$$

At the end of the first pass:

$$\texttt{temp\_keys} \; = \; \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & 010 & 110 & 100 & 001 & 101 & 011 & 111 \\ \end{array}}$$

$$\texttt{localOffsets} \; = \; \boxed{\begin{array}{c|c} 4 & 4 \end{array}}$$

At the end of the pass, we copy the result[3] of `temp_keys` into `keys`.

<u>Second pass</u>

The second pass now focuses on the second bit. The result is:[4]

$$\texttt{keys} = \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & 100 & 001 & 101 & 010 & 110 & 011 & 111 \\ \end{array}}$$

<u>Third (and last) pass</u>

$$\texttt{keys} = \boxed{\begin{array}{c|c|c|c|c|c|c|c} 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \\ \end{array}}$$

1.4 Parallel implementation

We first divide the array into blocks of size `blockSize` (here `blockSize` = 2).

$$\texttt{keys} = \underbrace{\boxed{\begin{array}{c|c} 001 & 101 \end{array}}}_{\text{block 0}} \underbrace{\boxed{\begin{array}{c|c} 011 & 000 \end{array}}}_{\text{block 1}} \underbrace{\boxed{\begin{array}{c|c} 010 & 111 \end{array}}}_{\text{block 2}} \underbrace{\boxed{\begin{array}{c|c} 110 & 100 \end{array}}}_{\text{block 3}}$$

Instead of creating a global histogram, we will create `numBlocks` local histograms using `computeBlockHistograms`:

$$\texttt{blockHistograms} = \underbrace{\boxed{\begin{array}{c|c} 0 & 2 \end{array}}}_{\text{block 0}} \underbrace{\boxed{\begin{array}{c|c} 1 & 1 \end{array}}}_{\text{block 1}} \underbrace{\boxed{\begin{array}{c|c} 1 & 1 \end{array}}}_{\text{block 2}} \underbrace{\boxed{\begin{array}{c|c} 2 & 0 \end{array}}}_{\text{block 3}}$$

---

[3]In fact, there is a way of avoiding this copy by storing alternatively the result in `keys` and `temp_keys`. If the number of passes is odd, there is a final copy from `temp_keys` to `keys`. This is sometimes called *ping-ponging.*

[4]Pay attention to the crucial role played by the stability property of Radix Sort!

We then combine these histograms into a global one using `reduceLocalHistoToGlobal` (notice that this produces the same result as the sequential version):

$$\texttt{globalHisto} = \boxed{4 \mid 4}$$

We scan this global histogram using `ScanGlobalHisto`:

$$\texttt{globalHistoExScan} = \boxed{0 \mid 4}$$

We then compute the offset for each of the local histograms using `computeBlockExScanFromGlobalHisto`. For this, we start by appending `globalHistoExScan` to the front of `blockHistograms` and discarding the final block as follows

$$\boxed{0 \mid 4 \mid 0 \mid 2 \mid 1 \mid 1 \mid 1 \mid 1}$$

We then scan this bucket-wise to get `blockExScan`.

$$\texttt{blockExScan} = \underbrace{\boxed{0 \mid 4}}_{\text{block 0}} \underbrace{\boxed{0 \mid 6}}_{\text{block 1}} \underbrace{\boxed{1 \mid 7}}_{\text{block 2}} \underbrace{\boxed{2 \mid 8}}_{\text{block 3}}$$

Each block in `blockExScan` represents the bucket-wise offsets in `temp_keys` for that block. The elements in `keys` should be written starting from this offset. As an example, the value 001 is the first value in block 0 and belongs to the second bucket (since its lsd is 1), it should therefore be written to position 4 in `temp_keys`. The value 101 follows the same pattern and must be written to the next position in `temp_keys`. In this case, that position is 5 (i.e., $4 + 1$). We populate the (partially) sorted array using `populateOutputFromBlockExScan`:

- block 0 will write at positions 4 and 5

- block 1 will write at positions 6 and 0

- block 2 will write at positions 1 and 7

- block 3 will write at positions 2 and 3