

# PARADIS: A Parallel Algorithm For In-Place Radix Sort

CME 323 Project

---

ROLLAND HE

# Radix Sort


---

- Non-comparative integer sorting algorithm
- Time complexity:  $O(kn)$ , where  $k$  is largest number of digits of any element in the list
- Looks at single digits at a time, making  $k$  passes through the data
- Places elements in buckets instead of comparing them

# Radix Sort

---

Runtime of Radix Sort vs Merge Sort



Dataset 1 – 1,000,000 random integers between 0 and 9,999

Dataset 2 – 1,000,000 random integers between 0 and  $9.99e15$

Dataset 3 – 100,000 random integers between 0 and  $9.99e15$

# PARADIS

---

- Introduced in 2014 by researchers Cho, Brand, Bordawekar, Finkler, Kulandaisamy, and Puri from IBM
- Parallel in-place radix sort algorithm
- Prior ideas:
  - Parallelizing the building of the histogram
- 2 new main ideas:
  - Speculative permutation/repair
  - Distribution-adaptive load balancing technique

# PARADIS - Analysis

---

- Time complexity –  $O\left(n\left(\frac{1}{p} + w\right)\right)$ , where  $p$  is the number of processors and  $w$  is the maximum fraction of elements that need to be repaired
  - Theoretical optimum –  $O\left(\frac{n}{p}\right)$
- Depth –  $O(\log n)$