

Algorithm Design and Analysis

Integer models of computation and integer sorting

Roadmap for today

- Breaking out of the comparison model, the *word-RAM*
- Learn about the *Counting Sort* algorithm
- Learn about the *Radix Sort* algorithm

Last Lecture: Sorting cannot be done faster than $\Omega(n \log n)$ ***in the comparison model***

Today: Sorting in $O(n)$ time ***for bounded integers in the word RAM model***

Formal model of computation

- We're leaving the comparison model today. We want to take advantage of integer inputs for more performance
- *Model (word-RAM):*
 - Unlimited constant-time addressable memory ("registers")
 - Each register can store a w -bit integer (a "word")
 - Reading/writing, arithmetic, logic, bitwise operations on a constant number of words takes constant time
 - With input size n , we need $w \geq \log n$ so that $2^w \geq n$ (in practice, $w = 64$)
- Assumption: w is large enough that all integers in the input to the problem fit in a single word

Implications of the word-RAM

- Adding two b -bit integers gives a $(b + 1)$ -bit integer
- Multiplying two b -bit integers gives a $2b$ -bit integer
- A constant number of these is okay since the result fits in a constant number of registers
- What if we **multiply** n w -bit integers? We get a $\Theta(nw)$ -bit answer! This **does not fit** in a single/constant number of registers!
- Such an algorithm would therefore take **more than** $\Theta(n)$ time

Real-life equivalent

```
int product = 1;  
for (int i = 0; i < n; i++)  
    product *= a[i];
```

```
product = 1  
for i in range(n):  
    product *= a[i]
```

- Too much addition/multiplication can quickly lead to **overflow**
- Python will represent large integers for you, but multiplying them **is not constant time**

- The word RAM model is just the theoretical equivalent of **watch out for overflow**. Something you should already be thinking about when designing algorithms

Do we *really* need to restrict to finite w ?

- Suppose we allow reading/writing/instructions on arbitrarily long integers
- This is usually called the unit-cost RAM (as opposed to the word-RAM)
- Can sort n arbitrarily large numbers in linear time [Paul, Simon]
- Pack n words into a single word of unlimited size and then 1 arithmetic operation on this big word performs n operations in parallel on our original words

ARTICLE FREE ACCESS

[Twitter](#) [LinkedIn](#) [Reddit](#) [Facebook](#) [Email](#)

A characterization of the class of functions computable in polynomial time on Random Access Machines

Then, we prove our main result: every problem in $\#P$ -SPACE can be solved in polynomial time by a RAM with the operations of sum, product, integer subtraction and integer division. The proof uses

Beating the comparison model

- As a warmup, consider the static searching problem

Problem (*Static search*) Given an array of elements a_1, a_2, \dots, a_n , with arbitrary preprocessing allowed for free, determine the index of a query element x if it exists

- What is a lower bound for this problem in the comparison model?

Static searching in the word RAM

- Suppose the array of elements are *integers* and we are in the word RAM model of computation
- **Preprocessing:** Build a *lookup table* S : If integer i was in position j in the input list a_1, \dots, a_n , then $S[i] = j$
- **Query:** To search for x , just look in position x and see if it's not empty

What is one problem with this approach?

The power of the word RAM

- The fundamental limitation of the comparison model is the fact that we can only have binary (YES / NO) decisions!
- The word RAM bestows upon us to lookup actual values in an array!
- A single instruction, e.g., lookup element i of an array of length n , can have many possible different outcomes!

Integer Sorting

Problem statement: Integer sorting

Problem (*Integer sorting*) We are given an array of elements a_1, a_2, \dots, a_n , each identified by a (not necessarily unique) **integer key** called ***key***(a_i).

Goal: output an array containing a permutation $a_{\pi_1}, a_{\pi_2}, \dots, a_{\pi_n}$ such that

$$\text{key}(a_{\pi_1}) \leq \text{key}(a_{\pi_2}) \leq \dots \leq \text{key}(a_{\pi_n})$$

<i>key</i>	<i>data</i>
key(a_1)	a_1
key(a_2)	a_2
key(a_3)	a_3

- Duplicate keys are allowed!
- Input contains arbitrary elements (not necessarily just integers) with ***integer keys***. Sorting must keep data + keys together

Stable Sorting

- Sorting is ***stable*** if relative order of duplicates is preserved
- If $\text{key}(a_i) = \text{key}(a_j)$ and $i < j$, then a_i comes before a_j in output


Name	Recitation
Dave	1
Alice	2
Ken	1
Eric	2
Carol	1

Sort
by
Name



Name	Recitation
Alice	2
Carol	1
Dave	1
Eric	2
Ken	1

Non-stable
Sort by
Recitation



Stable
Sort by
Recitation



Name	Recitation
Carol	1
Dave	1
Ken	1
Eric	2
Alice	2

Name	Recitation
Carol	1
Dave	1
Ken	1
Alice	2
Eric	2

Stable Sorting Continued

- Not all comparison-based sorting algorithms are stable
 - Quicksort is not stable (why?)
- Any comparison-based sorting algorithm can be made stable
- If $a_i = a_j$, then say $a_i < a_j$ if $i < j$, otherwise say $a_i > a_j$ if $i > j$

Sorting unique small integers



- We saw that we can beat the comparison model by taking advantage of indirect addressing (i.e., looking up in an array)
- ***Simpler problem:*** Suppose keys are guaranteed to be *unique* integers in $\{1, 2, \dots, n\}$

Algorithm:

- Create result array S of length n
- For each a_i , store $S[key(a_i)] = a_i$
- S is the sorted answer!

<i>key</i>	<i>data</i>
2	Cat
1	Dog
3	Tasselled Wobbegong



<i>key</i>	<i>data</i>
1	Dog
2	Cat
3	Tasselled Wobbegong

Sorting unique small integers

- Now let's increase the size of the keys. Suppose the input elements all have unique keys in range $\{0, 1, \dots, u - 1\}$ (the parameter u is called the ***universe*** of keys)

Algorithm:

- Create result array S of length u
- For each a_i , store $S[key(a_i)] = a_i$
- Filter out the empty elements of S
- S is the sorted answer!

Sorting small integers

- Now let's remove the assumption that the keys are unique. Suppose the input elements have (not necessarily unique) keys in the range $\{0, 1, \dots, u - 1\}$.

Algorithm (Counting Sort):

- Create a list for every possible key $\{0, 1, \dots, u - 1\}$
- For each a_i , append a_i to list at index $key(a_i)$
- Concatenate all the lists together
- Elements are sorted and Counting Sort is stable!

Counting Sort Example

Name	Recitation
Alice	2
Carol	1
Dave	1
Eric	2
Ken	1

Counting
Sort by
Recitation



Theorem: Counting Sort runs in $O(n + u)$ time
This is $O(n)$ time if $u = O(n)$

Side quest: Tuple sorting

Problem (*Tuple sorting*) Given an array of elements a_1, a_2, \dots, a_n , each identified by a **tuple of keys** (k_1, k_2, \dots, k_d) , sort the array **lexicographically** by the tuple. That is, the array is sorted by k_1 , with ties broken by k_2 , and ties on that broken by k_3 and so on!

2024, Feb, 16
2024, Feb, 18
2023, Feb, 06
2024, Jan, 16
2023, Jan, 19
2023, Jan, 17
2024, Feb, 06
2023, Feb, 07
2023, Feb, 19
2024, Feb, 15
2024, Jan, 19
2024, Jan, 18



2023, Jan, 17
2023, Jan, 19
2023, Feb, 06
2023, Feb, 07
2023, Feb, 19
2024, Jan, 16
2024, Jan, 18
2024, Jan, 19
2024, Feb, 06
2024, Feb, 15
2024, Feb, 16
2024, Feb, 18

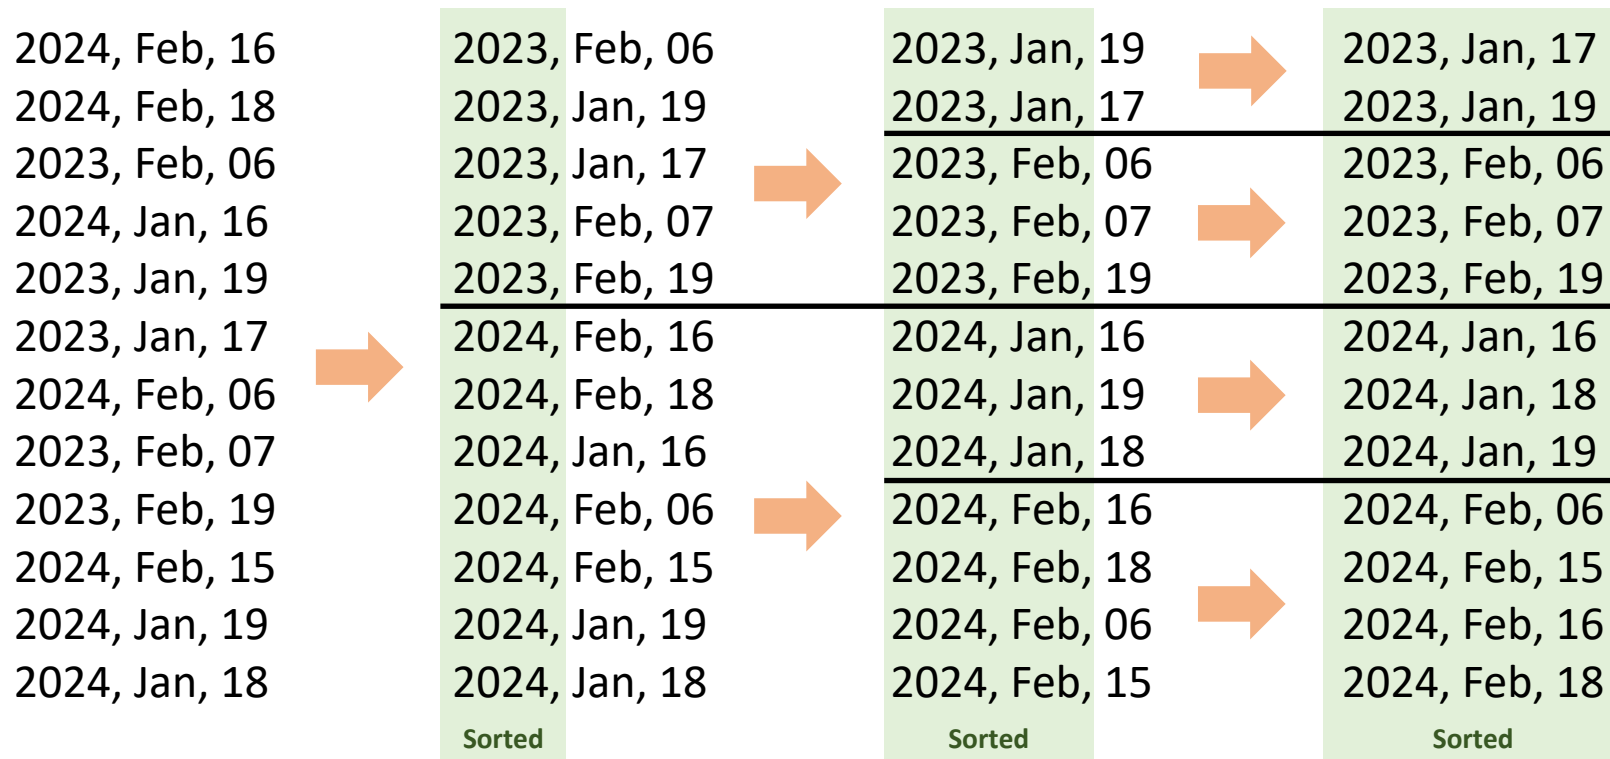
Algorithms for tuple sorting

Algorithm (Comparison tuple sort): Just use your favorite comparison-sorting algorithm (MergeSort, HeapSort, QuickSort, etc.) and compare tuples lexicographically

- **Cost:** $O(d n \log n)$ in the comparison model

Top-down tuple sorting

Algorithm (Top-down tuple sort): Sort by the first tuple element, then recursively sort the ties on the second tuple element and so on...



Bottom-up tuple sorting

Algorithm (Bottom-up tuple sort): **Stable sort** by the **last** tuple element, then the **second last**, and so on, finally sorting by the **first** tuple element

2024, Feb, 16
2024, Feb, 18
2023, Feb, 06
2024, Jan, 16
2023, Jan, 19
2023, Jan, 17
2024, Feb, 06
2023, Feb, 07
2023, Feb, 19
2024, Feb, 15
2024, Jan, 19
2024, Jan, 18



2023, Feb, 06
2024, Feb, 06
2023, Feb, 07
2024, Feb, 15
2024, Feb, 16
2024, Jan, 16
2023, Jan, 17
2024, Feb, 18
2024, Jan, 18
2023, Jan, 19
2023, Feb, 19
2024, Jan, 19

Sorted



2024, Jan, 16
2023, Jan, 17
2024, Jan, 18
2023, Jan, 19
2024, Jan, 19
2023, Feb, 06
2024, Feb, 06
2023, Feb, 07
2024, Feb, 15
2024, Feb, 16
2024, Feb, 18
2023, Feb, 19

Sorted



2023, Jan, 17
2023, Jan, 19
2023, Feb, 06
2023, Feb, 07
2023, Feb, 19
2024, Jan, 16
2024, Jan, 18
2024, Jan, 19
2024, Feb, 06
2024, Feb, 15
2024, Feb, 16
2024, Feb, 18

Sorted

Sorting bigger integers


- Counting sort runs in linear ($O(n)$) time for $u = O(n)$
- We want to sort in linear time for bigger values of u
- **Idea:** *Use tuple sort* to sort integer keys

Question: Can we represent a big integer as a tuple of small integers such that tuple sorting them gives the right answer?

Answer: Just use their *digits*! (Small integers so Counting Sort works)

Bottom-up (LSD) Radix Sort

Algorithm (LSD Radix Sort): Counting sort by the **last** digit, then the **second last**, and so on, finally sorting by the **first** digit.



Theorem: Radix Sort runs in $O((n + b) \log_b u)$ time using base- b

Optimal choice of b ?

- How do we optimize $O((n + b) \log_b u)$?
- Bigger base \Rightarrow fewer iterations (but also slower Counting Sort)

Optimal base: $b =$

Running time:

Theorem: Radix Sort can sort keys in $\{0, 1, \dots, O(n^c)\}$ in $O(n)$ time!

Summary of Radix and Counting Sort

Given n input elements with integer keys in $\{0, 1, \dots, u - 1\}$,

- **Counting Sort** runs in $O(n + u)$ time.
 - This is linear time whenever $u = O(n)$, i.e., linear-sized keys
- **Radix Sort** runs in $O(n \log_n u)$ time.
 - This is linear time whenever $u = O(n^c)$, i.e., polynomial-sized keys!

Fun fact (Integer sorting is still an open problem): We don't know whether there exists an algorithm that can sort integers of any size in linear time. The best discovered algorithms take $O(n \log \log n)$ time (deterministic) or $O(n \sqrt{\log \log n})$ expected time. No known lower bound proves that linear-time integer sorting is impossible, but we don't know!

Applications and Extensions

- How to solve this: given n integers a_1, a_2, \dots, a_n in $\{1, 2, 3, \dots, n^2\}$, find an $i \neq j$ for which $|a_i - a_j|$ is minimized
 - Use Radix Sort, then walk through the sorted order and maintain a smallest consecutive difference
- Most Significant Digit (MSD) Radix Sort sorts the MSD first then recurses in each bucket of items with the same MSD
- What might go wrong with MSD Radix Sort?
 - CountingSort applied in each recursive call, takes $O(u)$ time regardless of number of items in recursive call
 - If $\Theta(n)$ recursive calls each of $O(1)$ items, then $O(n \cdot u)$ time if done naively