# CS60020: Foundations of <br> Algorithm Design and Machine Learning <br> <br> Sourangshu Bhattacharya 

 <br> <br> Sourangshu Bhattacharya}

## How fast can we sort?

All the sorting algorithms we have seen so far are comparison sorts: only use comparisons to determine the relative order of elements.

- E.g., insertion sort, merge sort, quicksort, heapsort.
The best worst-case running time that we've seen for comparison sorting is $O(n \lg n)$.


## Is O(nlgn) the best we can do?

Decision trees can help us answer this question.

## Decision-tree example

## Sort $\left\langle a_{1}, a_{2}, \ldots, a_{n}\right\rangle$



Each internal node is labeled $i: j$ for $i, j \in\{1,2, \ldots, n\}$.

- The left subtree shows subsequent comparisons if $a_{i} \leq a_{j}$.
- The right subtree shows subsequent comparisons if $a_{i} \geq a_{j}$.


## Decision-tree example

## Sort $\left\langle a_{1}, a_{2}, a_{3}\right\rangle$ <br> $=\langle 9,4,6\rangle$ :



Each internal node is labeled $i: j$ for $i, j \in\{1,2, \ldots, n\}$.

- The left subtree shows subsequent comparisons if $a_{i} \leq a_{j}$.
- The right subtree shows subsequent comparisons if $a_{i} \geq a_{j}$.


## Decision-tree example

$$
\begin{aligned}
& \text { Sort }\left\langle a_{1}, a_{2}, a_{3}\right\rangle \\
& =\langle 9,4,6\rangle \text { : }
\end{aligned}
$$



Each internal node is labeled $i: j$ for $i, j \in\{1,2, \ldots, n\}$.

- The left subtree shows subsequent comparisons if $a_{i} \leq a_{j}$.
- The right subtree shows subsequent comparisons if $a_{i} \geq a_{j}$.


## Decision-tree example

## Sort $\left\langle a_{1}, a_{2}, a_{3}\right\rangle$ $=\langle 9,4,6\rangle$ :



Each internal node is labeled $i: j$ for $i, j \in\{1,2, \ldots, n\}$.

- The left subtree shows subsequent comparisons if $a_{i} \leq a_{j}$.
- The right subtree shows subsequent comparisons if $a_{i} \geq a_{j}$.


## Decision-tree example

Sort $\left\langle a_{1}, a_{2}, a_{3}\right\rangle$ $=\langle 9,4,6\rangle$ :


Each leaf contains a permutation $\langle\pi(1), \pi(2), \ldots, \pi(n)\rangle$ to indicate that the ordering $a_{\pi(1)} \leq a_{\pi(2)} \leq \mathrm{L} \leq a_{\pi(\mathrm{n})}$ has been established.

## Decision-tree model

A decision tree can model the execution of any comparison sort:

- One tree for each input size $n$.
- View the algorithm as splitting whenever it compares two elements.
- The tree contains the comparisons along all possible instruction traces.
- The running time of the algorithm $=$ the length of the path taken.
- Worst-case running time $=$ height of tree.


## Lower bound for decision- tree

## sorting

Theorem. Any decision tree that can sort $n$ elements must have height $\Omega(n \lg n)$.

Proof. The tree must contain $\geq n$ ! leaves, since there are $n$ ! possible permutations. A height- $h$ binary tree has $\leq 2^{h}$ leaves. Thus, $n!\leq 2^{h}$.
$\therefore h \geq \lg (n!)$
$\geq \lg \left((n / e)^{n}\right)$
( lg is mono. increasing)
$=n \lg n-n \lg e$
$=\Omega(n \lg n)$.
(Stirling's formula)

## Lower bound for comparison sorting

## Sorting in linear time

## Counting sort: No comparisons between elements.

- Input: $A[1$. . n], where $A[j] \in\{1,2, \ldots, k\}$.
- Output: $B[1$. . n], sorted.
- Auxiliary storage: $C[1 . . k]$.


## Counting sort

for $i \leftarrow 1$ to $k$

## do $C[i] \leftarrow 0$

for $j \leftarrow 1$ to $n$

$$
\text { do } C[A[j]] \leftarrow C[A[j]]+1 \quad \triangleleft C[i]=\mid\{\text { key }=i\} \mid
$$

for $i \leftarrow 2$ to $k$
do $C[i] \leftarrow C[i]+C[i-1]$
$\triangleleft C[i]=|\{\mathrm{key} \leq i\}|$
for $j \leftarrow n$ downto 1
do $B[C[A[j]]] \leftarrow \mathrm{A}[j]$

$$
C[A[j]] \leftarrow C[A[j]]-1
$$

## Counting-sort example



$\mathbf{f o r} i \leftarrow 1$ to $k$

## do $C[i] \leftarrow 0$

## Loop 2


$\boldsymbol{f o r} j \leftarrow 1$ to $n$
$\mathbf{d o} C[A[j]] \leftarrow C[A[j]]+1 \quad \triangleleft C[i]=|\{\mathrm{key}=i\}|$

## Loop 2


$\boldsymbol{f o r} j \leftarrow 1$ to $n$
do $C[A[j]] \leftarrow C[A[j]]+1 \quad \triangleleft C[i]=\mid\{$ key $=i\} \mid$

## Loop 2


$\boldsymbol{f o r} j \leftarrow 1$ to $n$
do $C[A[j]] \leftarrow C[A[j]]+1 \quad \triangleleft C[i]=\mid\{$ key $=i\} \mid$

## Loop 2


for $j \leftarrow 1$ to $n$
do $C[A[j]] \leftarrow C[A[j]]+1 \quad \triangleleft C[i]=\mid\{$ key $=i\} \mid$

## Loop 2

| 4 |  | 4 | 3 |  | 1 | 0 |  | 2 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

$B$ :

|  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |

$\boldsymbol{f o r} j \leftarrow 1$ to $n$
$\boldsymbol{d o} C[A[j]] \leftarrow C[A[j]]+1 \quad \triangleleft C[i]=\mid\{$ key $=i\} \mid$

$A:$|  | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: |
| 4 | 1 | 3 | 4 | 3 |


$B$ :



for $i \leftarrow 2$ to $k$

$$
\boldsymbol{d o} C[i] \leftarrow C[i]+C[i-1] \quad \triangleleft C[i]=\mid\{\text { key } \leq i\} \mid
$$

$A:$|  | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: |
| 4 | 1 | 3 | 4 | 3 |


$B$ :

|  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |


$\mathbf{f o r} i \leftarrow 2$ to $k$

$$
\boldsymbol{d o} C[i] \leftarrow C[i]+C[i-1] \quad \triangleleft C[i]=\mid\{\text { key } \leq i\} \mid
$$

$A:$|  | 1 | 2 | 3 | 4 |
| :---: | :---: | :---: | :---: | :---: |
| 4 | 1 | 3 | 4 | 3 |


$\mathbf{f o r} i \leftarrow 2 \boldsymbol{t o} k$

$$
\boldsymbol{d o} C[i] \leftarrow C[i]+C[i-1] \quad \triangleleft C[i]=\mid\{\text { key } \leq i\} \mid
$$

## Loop 4


for $j \leftarrow n$ downto 1

$$
\begin{aligned}
& \text { do } B[C[A[j]]] \leftarrow \mathrm{A}[j] \\
& \quad C[A[j]] \leftarrow C[A[j]]-1
\end{aligned}
$$

## Loop 4


for $j \leftarrow n$ downto 1

$$
\begin{aligned}
& \text { do } B[C[A[j]]] \leftarrow \mathrm{A}[j] \\
& \quad C[A[j]] \leftarrow C[A[j]]-1
\end{aligned}
$$


for $j \leftarrow n$ downto 1

$$
\begin{aligned}
& \text { do } B[C[A[j]]] \leftarrow \mathrm{A}[j] \\
& \quad C[A[j]] \leftarrow C[A[j]]-1
\end{aligned}
$$


for $j \leftarrow n$ downto 1

$$
\begin{aligned}
& \text { do } B[C[A[j]]] \leftarrow \mathrm{A}[j] \\
& \quad C[A[j]] \leftarrow C[A[j]]-1
\end{aligned}
$$

## Loop 4


for $j \leftarrow n$ downto 1

$$
\begin{aligned}
& \text { do } B[C[A[j]]] \leftarrow \mathrm{A}[j] \\
& \quad C[A[j]] \leftarrow C[A[j]]-1
\end{aligned}
$$

## Analysis

$\Theta(k) \quad\left\{\begin{array}{l}\text { for } i \leftarrow 1 \text { to } k \\ \text { do } C[i] \leftarrow 0\end{array}\right.$
$\Theta(n) \quad\left\{\begin{array}{l}\text { for } j \leftarrow 1 \text { to } n\end{array}\right.$ do $C[A[j]] \leftarrow C[A[j]]+1$
$\Theta(k) \quad\left\{\begin{array}{l}\text { for } i \leftarrow 2 \text { to } k \\ \text { do } C[i] \leftarrow C[i]+C[i-1]\end{array}\right.$
$\Theta(n)\left\{\begin{array}{l}\text { for } j \leftarrow n \text { downto } 1 \\ \quad \text { do } B[C[A[j]]] \leftarrow \mathrm{A}[j]\end{array}\right.$ $C[A[j]] \leftarrow C[A[j]]-1$
$\Theta(n+k)$

## Running time

If $k=O(n)$, then counting sort takes $\Theta(n)$ time.

- But, sorting takes $\Omega(n \lg n)$ time!
- Where's the fallacy?

Answer:

- Comparison sorting takes $\Omega(n \lg n)$ time.
- Counting sort is not a comparison sort.
- In fact, not a single comparison between elements occurs!


## Stable sorting

Counting sort is a stable sort: it preserves the input order among equal elements.


Exercise: What other sorts have this property?

## Radix sort

- Origin: Herman Hollerith's card-sorting machine for the 1890 U.S. Census. (See Appendix (0)
- Digit-by-digit sort.
- Hollerith's original (bad) idea: sort on most-significant digit first.
- Good idea: Sort on least-significant digit first with auxiliary stable sort.


## Operation of radix sort

| 329 | 720 | 720 | 329 |
| ---: | ---: | ---: | ---: | ---: |
| 457 | 355 | 329 | 355 |
| 657 | 436 | 436 | 436 |
| 839 | 457 | 839 | 457 |
| 436 | 657 | 355 | 657 |
| 720 | 329 | 457 | 720 |
| 355 | 839 | 657 | 839 |
|  |  |  |  |
|  |  |  |  |

## Correctness of radix sort

Induction on digit position

- Assume that the numbers are sorted by their low-order $t-1$ digits.
- Sort on digit $t$

| 720 | 329 |
| :--- | :--- | :--- |
| 329 | 355 |
| 436 | 436 |
| 839 | 457 |
| 355 | 657 |
| 457 | 720 |
| 657 | 839 |

## Correctness of radix sort

Induction on digit position

- Assume that the numbers are sorted by their low-order $t-1$ digits.
- Sort on digit $t$
- Two numbers that differ in digit $t$ are correctly sorted.



## Correctness of radix sort

Induction on digit position

- Assume that the numbers are sorted by their low-order $t-1$ digits.
- Sort on digit $t$
- Two numbers that differ in digit $t$ are correctly sorted.
- Two numbers equal in digit $t$ are put in the same order as the input $\Rightarrow$ correct order.



## Analysis of radix sort

- Assume counting sort is the auxiliary stable sort.
- Sort $n$ computer words of $b$ bits each.
- Each word can be viewed as having $b / r$ base- $2^{r}$ digits.
Example: 32-bit word

$r=8 \Rightarrow b / r=4$ passes of counting sort on base- $2^{8}$ digits; or $r=16 \Rightarrow b / r=2$ passes of counting sort on base- $2{ }^{16}$ digits.


## How many passes should we make?

## Analysis (continued)

Recall: Counting sort takes $\Theta(n+k)$ time to sort $n$ numbers in the range from 0 to $k-1$.
If each $b$-bit word is broken into $r$-bit pieces, each pass of counting sort takes $\Theta\left(n+2^{r}\right)$ time. Since there are $b / r$ passes, we have

$$
T(n, b)=\Theta\left(\frac{b}{r}\left(n+2^{r}\right)\right) .
$$

Choose $r$ to minimize $T(n, b)$ :

- Increasing $r$ means fewer passes, but as $r \gg \lg n$, the time grows exponentially.


## Choosing $r$

$$
T(n, b)=\Theta\left(\frac{b}{r}\left(n+2^{r}\right)\right) .
$$

Minimize $T(n, b)$ by differentiating and setting to 0 .
Or, just observe that we don't want $2^{r} \gg n$, and there's no harm asymptotically in choosing $r$ as large as possible subject to this constraint.
Choosing $r=\lg n$ implies $T(n, b)=\Theta(b n / \lg n)$.

- For numbers in the range from 0 to $n^{d}-1$, we have $b=d \lg n \Rightarrow$ radix sort runs in $\Theta(d n)$ time.


## Conclusions

In practice, radix sort is fast for large inputs, as well as simple to code and maintain.
Example (32-bit numbers):

- At most 3 passes when sorting $\geq 2000$ numbers.
- Merge sort and quicksort do at least $\lg 2000=11$ passes.
Downside: Unlike quicksort, radix sort displays little locality of reference, and thus a well-tuned quicksort fares better on modern processors, which feature steep memory hierarchies.

