CS 5633 -- Spring 2005

INTRODUCTION TO
ALGORITHMS
SECOND EDITION

THOMAS H. CORMEN
CHARLES E. LEISERSON
RONALD L. RIVEST
CLIFFORD STEIN

Sorting

Carola Wenk

Slides courtesy of Charles Leiserson with small changes by Carola Wenk

---

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 $O(n \log n)$.

Is O(n log n) the best we can do?

Decision trees can help us answer this question.

Decision-tree example

Sort $\langle a_1, a_2, \ldots, a_n \rangle$

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 $\langle a_1, a_2, a_3 \rangle$
$= \langle\, 9, 4, 6\, \rangle$:

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 $\langle a_1, a_2, a_3 \rangle$
$= \langle\, 9, 4, 6\, \rangle$:

Each internal node is labeled $i{:}j$ for $i, j \in \{1, 2, \ldots, n\}$.
- The left subtree shows subsequent comparisons if $a_i \le a_j$.
- The right subtree shows subsequent comparisons if $a_i \ge a_j$.

Decision-tree example

Sort $\langle a_1, a_2, a_3 \rangle$
$= \langle\, 9, 4, 6\, \rangle$:

Each internal node is labeled $i{:}j$ for $i, j \in \{1, 2, \ldots, n\}$.
- The left subtree shows subsequent comparisons if $a_i \le a_j$.
- The right subtree shows subsequent comparisons if $a_i \ge a_j$.

Decision-tree example

Sort $\langle a_1, a_2, a_3 \rangle$
$= \langle\, 9, 4, 6\, \rangle$:

Each internal node is labeled $i{:}j$ for $i, j \in \{1, 2, \ldots, n\}$.
- The left subtree shows subsequent comparisons if $a_i \le a_j$.
- The right subtree shows subsequent comparisons if $a_i \ge a_j$.

Decision-tree example

Sort $\langle a_1, a_2, a_3 \rangle$
$= \langle\, 9, 4, 6\, \rangle$:

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

Decision-tree model

A decision tree can model the execution of any comparison sorting algorithm:

- 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 $\Omega(n \log n)$.

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, $n! \leq 2^h$.

$$\therefore h \geq \log(n!) \quad\quad (\log \text{ is mono. increasing})$$
$$\geq \log((n/e)^n) \quad\quad (\text{Stirling's formula})$$
$$= n \log n - n \log e$$
$$= \Omega(n \log n). \quad \square$$

Lower bound for comparison sorting

Corollary. Heapsort and merge sort are asymptotically optimal comparison sorting algorithms.

Sorting in linear time

Counting sort: No comparisons between elements.

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

Counting sort

for $i \leftarrow 1$ **to** k
   do $C[i] \leftarrow 0$
for $j \leftarrow 1$ **to** n
   do $C[A[j]] \leftarrow C[A[j]] + 1$ ▷ $C[i] = |\{key = i\}|$
for $i \leftarrow 2$ **to** k
   do $C[i] \leftarrow C[i] + C[i{-}1]$ ▷ $C[i] = |\{key \leq i\}|$
for $j \leftarrow n$ **downto** 1
   do $B[C[A[j]]] \leftarrow A[j]$
     $C[A[j]] \leftarrow C[A[j]] - 1$

Counting-sort example

Loop 1

for $i \leftarrow 1$ **to** k
   **do** $C[i] \leftarrow 0$

Loop 2

for $j \leftarrow 1$ **to** n
   do $C[A[j]] \leftarrow C[A[j]] + 1$ ▷ $C[i] = |\{key = i\}|$

Loop 2

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	0	0	1

B:					

for $j \leftarrow 1$ **to** n
    do $C[A[j]] \leftarrow C[A[j]] + 1$ ▷ $C[i] = |\{key = i\}|$

Loop 2

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	0	1	1

B:					

for $j \leftarrow 1$ **to** n
    do $C[A[j]] \leftarrow C[A[j]] + 1$ ▷ $C[i] = |\{key = i\}|$

Loop 2

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	0	1	2

B:					

for $j \leftarrow 1$ **to** n
    do $C[A[j]] \leftarrow C[A[j]] + 1$ ▷ $C[i] = |\{key = i\}|$

Loop 2

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	0	2	2

B:					

for $j \leftarrow 1$ **to** n
    do $C[A[j]] \leftarrow C[A[j]] + 1$ ▷ $C[i] = |\{key = i\}|$

Loop 3

A: | 4 | 1 | 3 | 4 | 3 |

C: | 1 | 0 | 2 | 2 |

B: | | | | | |

C′: | 1 | 1 | 2 | 2 |

for $i \leftarrow 2$ **to** k
   do $C[i] \leftarrow C[i] + C[i-1]$ ▷ $C[i] = |\{\text{key} \leq i\}|$

Loop 3

A: | 4 | 1 | 3 | 4 | 3 |

C: | 1 | 0 | 2 | 2 |

B: | | | | | |

C′: | 1 | 1 | 3 | 2 |

for $i \leftarrow 2$ **to** k
   do $C[i] \leftarrow C[i] + C[i-1]$ ▷ $C[i] = |\{\text{key} \leq i\}|$

Loop 3

A: | 4 | 1 | 3 | 4 | 3 |

C: | 1 | 0 | 2 | 2 |

B: | | | | | |

C′: | 1 | 1 | 3 | 5 |

for $i \leftarrow 2$ **to** k
   do $C[i] \leftarrow C[i] + C[i-1]$ ▷ $C[i] = |\{\text{key} \leq i\}|$

Loop 4

A: | 4 | 1 | 3 | 4 | 3 |

C: | 1 | 1 | 3 | 5 |

B: | | 3 | | | |

C′: | 1 | 1 | 3 | 5 |

for $j \leftarrow n$ **downto** 1
   do $B[C[A[j]]] \leftarrow A[j]$
      $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	1	3	5

	1	2	3	4	5
B:			3		

	1	2	3	4
C':	1	1	2	5

for $j \leftarrow n$ **downto** 1
   do $B[C[A[j]]] \leftarrow A[j]$
     $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	1	2	5

	1	2	3	4	5
B:			3		4

	1	2	3	4
C':	1	1	2	5

for $j \leftarrow n$ **downto** 1
   do $B[C[A[j]]] \leftarrow A[j]$
     $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	1	2	5

	1	2	3	4	5
B:			3		4

	1	2	3	4
C':	1	1	2	4

for $j \leftarrow n$ **downto** 1
   do $B[C[A[j]]] \leftarrow A[j]$
     $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	1	1	2	4

	1	2	3	4	5
B:		3	3		4

	1	2	3	4
C':	1	1	2	4

for $j \leftarrow n$ **downto** 1
   do $B[C[A[j]]] \leftarrow A[j]$
     $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

A (1–5): 4 1 3 4 3 C (1–4): 1 1 2 4

B: 3 3 4 C': 1 1 1 4

for $j \leftarrow n$ **downto** 1
  do $B[C[A[j]]] \leftarrow A[j]$
    $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

A (1–5): 4 1 3 4 3 C (1–4): 1 1 1 4

B: 1 3 3 4 C': 1 1 1 4

for $j \leftarrow n$ **downto** 1
  do $B[C[A[j]]] \leftarrow A[j]$
    $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

A (1–5): 4 1 3 4 3 C (1–4): 1 1 1 4

B: 1 3 3 4 C': 0 1 1 4

for $j \leftarrow n$ **downto** 1
  do $B[C[A[j]]] \leftarrow A[j]$
    $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

A (1–5): 4 1 3 4 3 C (1–4): 0 1 1 4

B: 1 3 3 4 4 C': 0 1 1 4

for $j \leftarrow n$ **downto** 1
  do $B[C[A[j]]] \leftarrow A[j]$
    $C[A[j]] \leftarrow C[A[j]] - 1$

Loop 4

	1	2	3	4	5
A:	4	1	3	4	3

	1	2	3	4
C:	0	1	1	4

	1	2	3	4	5
B:	1	3	3	4	4

	1	2	3	4
C':	0	1	1	3

for $j \leftarrow n$ **downto** 1
    do $B[C[A[j]]] \leftarrow A[j]$
        $C[A[j]] \leftarrow C[A[j]] - 1$

Analysis

$\Theta(k)$ $\left\{\begin{array}{l} \textbf{for } i \leftarrow 1 \textbf{ to } k \\ \quad \textbf{do } C[i] \leftarrow 0 \end{array}\right.$

$\Theta(n)$ $\left\{\begin{array}{l} \textbf{for } j \leftarrow 1 \textbf{ to } n \\ \quad \textbf{do } C[A[j]] \leftarrow C[A[j]] + 1 \end{array}\right.$

$\Theta(k)$ $\left\{\begin{array}{l} \textbf{for } i \leftarrow 2 \textbf{ to } k \\ \quad \textbf{do } C[i] \leftarrow C[i] + C[i-1] \end{array}\right.$

$\Theta(n)$ $\left\{\begin{array}{l} \textbf{for } j \leftarrow n \textbf{ downto } 1 \\ \quad \textbf{do } B[C[A[j]]] \leftarrow A[j] \\ \qquad C[A[j]] \leftarrow C[A[j]] - 1 \end{array}\right.$

$\Theta(n + k)$

Running time

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

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

Answer:

- ***Comparison sorting*** takes $\Omega(n \log n)$ 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.

A:	4	1	3	4	3

|					
B:	1	3	3	4	4

Exercise: What other sorts have this property?

Radix sort

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

Operation of radix sort

3 2 9	7 2 0	7 2 0	3 2 9
4 5 7	3 5 5	3 2 9	3 5 5
6 5 7	4 3 6	4 3 6	4 3 6
8 3 9	4 5 7	8 3 9	4 5 7
4 3 6	6 5 7	3 5 5	6 5 7
7 2 0	3 2 9	4 5 7	7 2 0
3 5 5	8 3 9	6 5 7	8 3 9

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

7 2 0	3 2 9
3 2 9	3 5 5
4 3 6	4 3 6
8 3 9	4 5 7
3 5 5	6 5 7
4 5 7	7 2 0
6 5 7	8 3 9

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.

7 2 0	3 2 9
3 2 9	3 5 5
4 3 6	4 3 6
8 3 9	4 5 7
3 5 5	6 5 7
4 5 7	7 2 0
6 5 7	8 3 9

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.

7 2 0		3 2 9
3 2 9		3 5 5
4 3 6	→	4 3 6
8 3 9		4 5 7
3 5 5		6 5 7
4 5 7		7 2 0
6 5 7		8 3 9

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

8	8	8	8

$r = 8 \Rightarrow b/r = 4$ 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 $\Theta(n + 2^r)$ 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 $r \gg \log n$, 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 $r = \log n$ implies $T(n, b) = \Theta(bn/\log n)$.

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

Radix Sort

- Assume counting sort is the auxiliary stable sort.
- Sort n computer words of b bits each.

The runtime of radix sort is: $T(n, b) = \Theta(bn/\log n)$.

- Example:
  For numbers in the range from 0 to $n^d - 1$, we have $b = d \log n \Rightarrow$ radix sort runs in $\Theta(d\,n)$ time.
- Notice that counting sort runs in $O(n+k)$ time, where all numbers are in the range 1 through k.

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 ≥ 2000 numbers.
- Merge sort and quicksort do at least $\lceil \log 2000 \rceil = 11$ 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.

Appendix: Punched-card technology

- Herman Hollerith (1860-1929)
- Punched cards
- Hollerith's tabulating system
- Operation of the sorter
- Origin of radix sort
- "Modern" IBM card

Return to last slide viewed.


Herman Hollerith (1860-1929)

- The 1880 U.S. Census took almost 10 years to process.
- While a lecturer at MIT, Hollerith prototyped punched-card technology.
- His machines, including a "card sorter," allowed the 1890 census total to be reported in 6 weeks.
- He founded the Tabulating Machine Company in 1911, which merged with other companies in 1924 to form International Business Machines.

Punched cards

- Punched card = data record.
- Hole = value.
- Algorithm = machine + human operator.
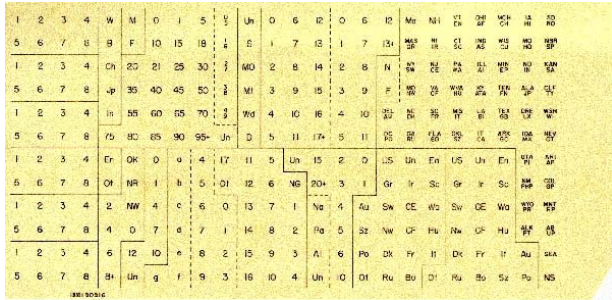
Replica of punch card from the 1900 U.S. census. [Howells 2000]

Hollerith's tabulating system

Figure from [Howells 2000].

- Pantograph card punch
- Hand-press reader
- Dial counters
- Sorting box


Operation of the sorter

- An operator inserts a card into the press.
- Pins on the press reach through the punched holes to make electrical contact with mercury-filled cups beneath the card.
- Whenever a particular digit value is punched, the lid of the corresponding sorting bin lifts.
- The operator deposits the card into the bin and closes the lid.

Hollerith Tabulator, Pantograph, Press, and Sorter

- When all cards have been processed, the front panel is opened, and the cards are collected in order, yielding one pass of a stable sort.

Origin of radix sort

Hollerith's original 1889 patent alludes to a most-significant-digit-first radix sort:

"The most complicated combinations can readily be counted with comparatively few counters or relays by first assorting the cards according to the first items entering into the combinations, then reassorting each group according to the second item entering into the combination, and so on, and finally counting on a few counters the last item of the combination for each group of cards."

Least-significant-digit-first radix sort seems to be a folk invention originated by machine operators.

"Modern" IBM card

• One character per column.

Produced by the [WWW Virtual Punch-Card Server](#).

So, that's why text windows have 80 columns!