

CMPS 2200 – Fall 2017

Sorting

Carola Wenk

Slides courtesy of Charles Leiserson with changes
and additions 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.

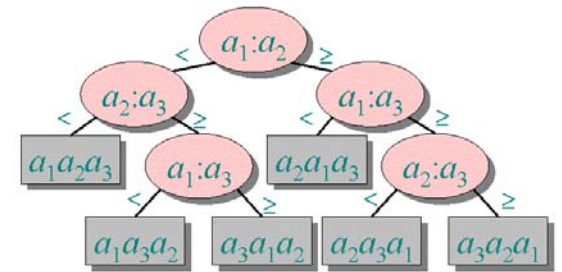
- *E.g.*, insertion sort, merge sort, heapsort.

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 model



A decision tree models the execution of any comparison sorting algorithm:

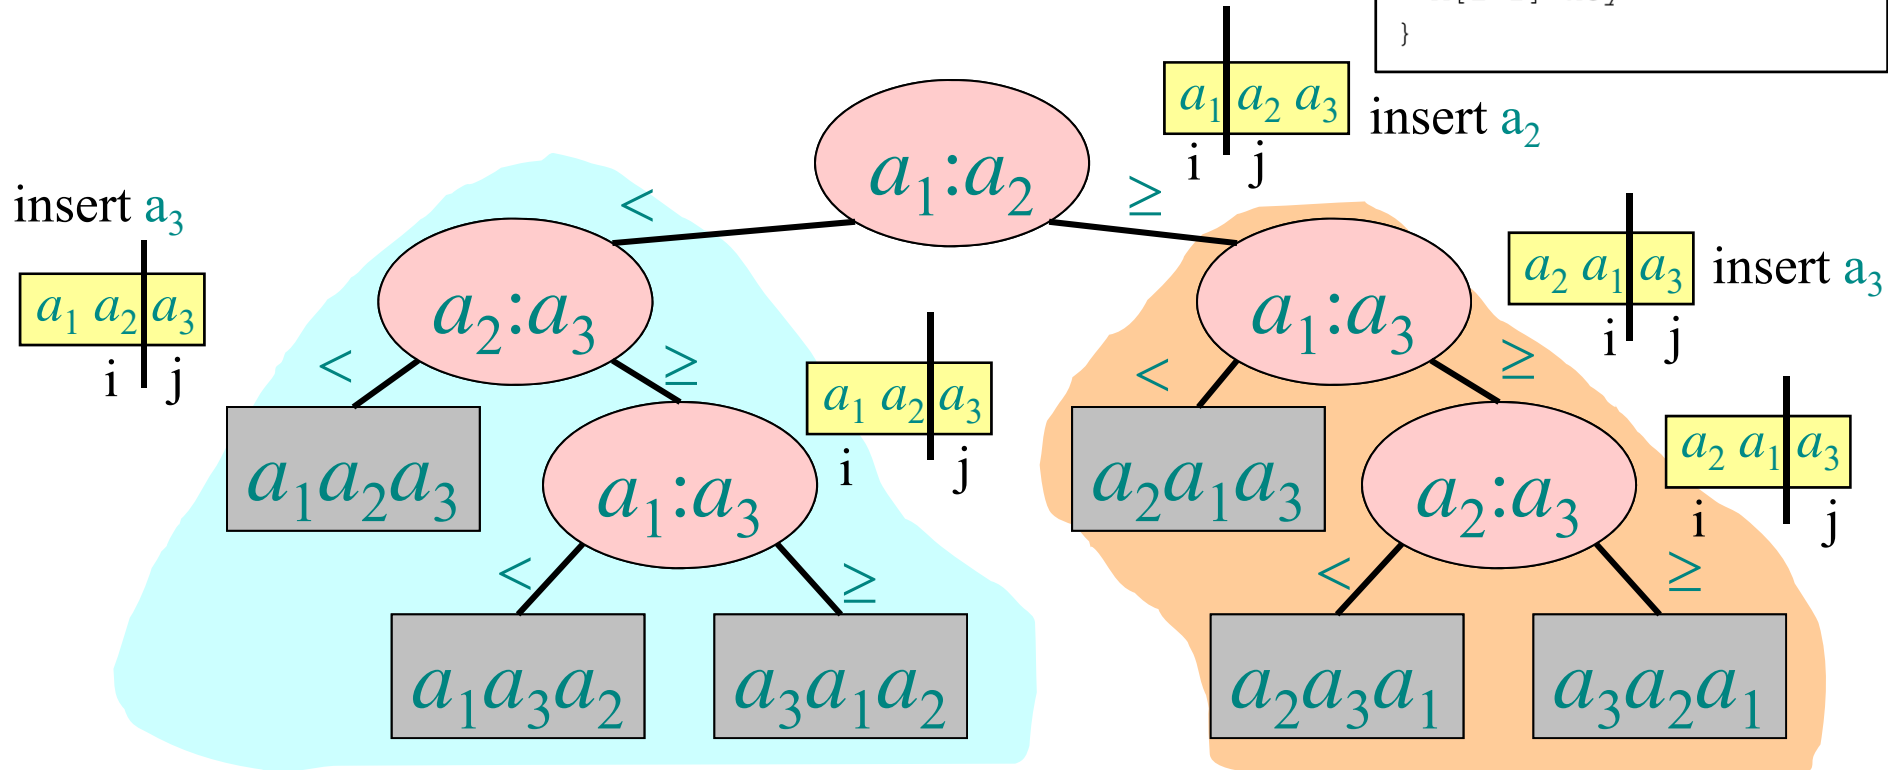
- One tree per input size n .
- The tree contains **all** possible comparisons (= if-branches) that could be executed for **any** input of size n .
- The tree contains **all** comparisons along **all** possible instruction traces (= control flows) for **all** inputs of size n .
- For one input, only one path to a leaf is executed.
- Running time = length of the path taken.
- Worst-case running time = height of tree.

Decision-tree for insertion sort

sort: Sort $\langle a_1, a_2, a_3 \rangle$

```

for j=2 to n {
  key = A[j]
  // Insert A[j] into
  // sorted sequence A[1..j-1]
  i=j-1
  while(i>0 && A[i]>key) {
    A[i+1]=A[i]
    i--
  }
  A[i+1]=key
}
    
```

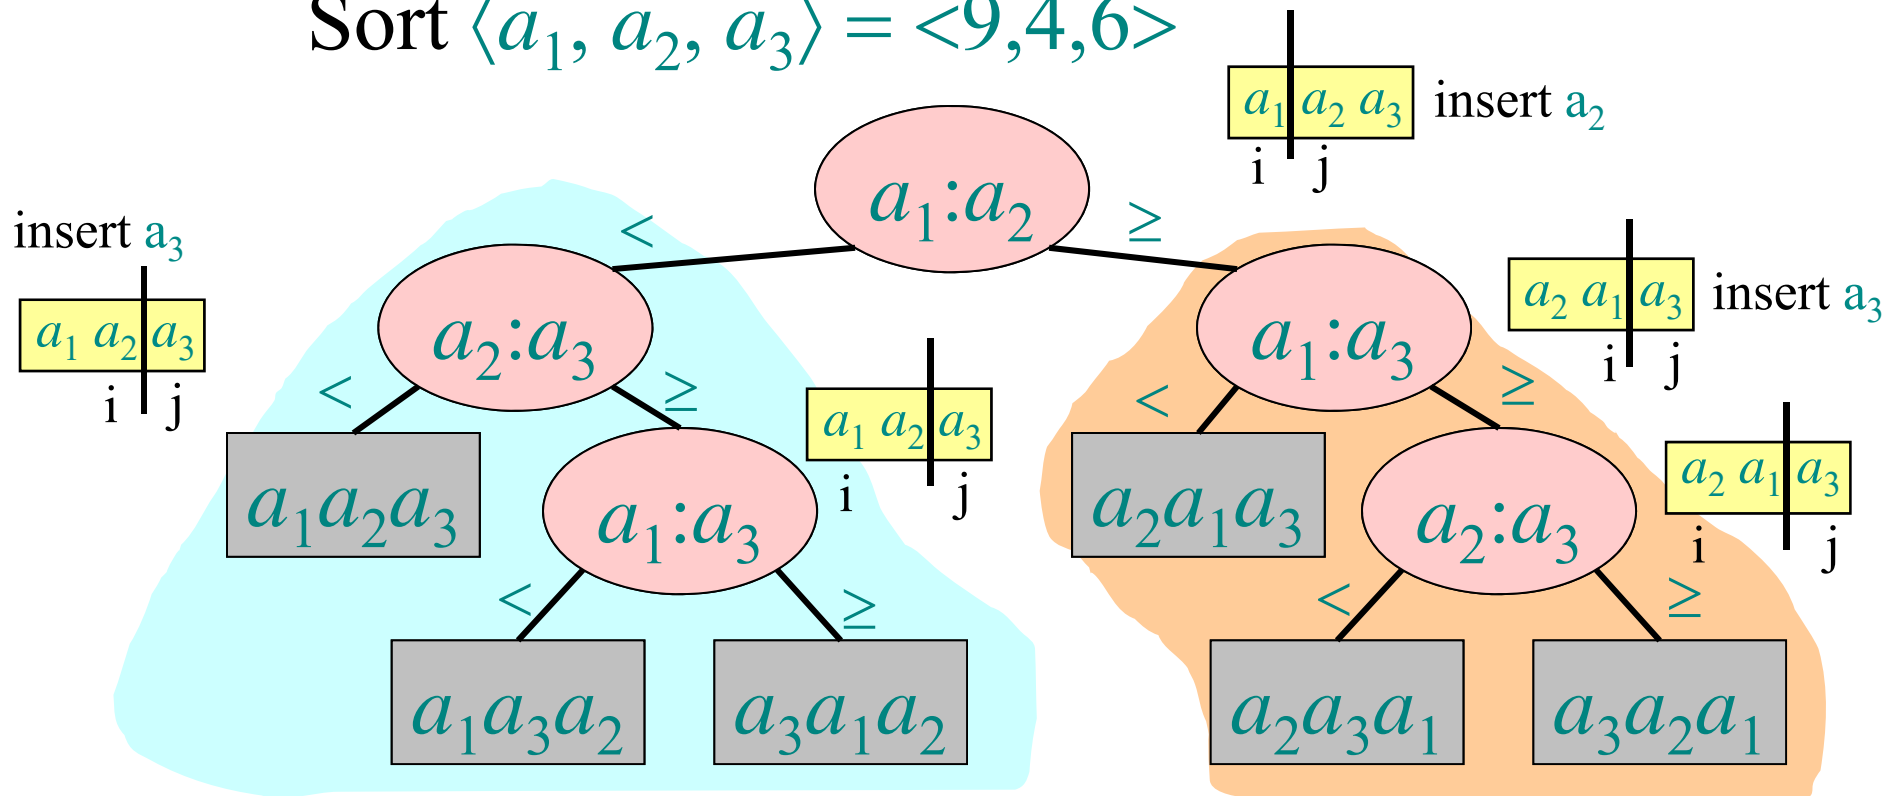


Each internal node is labeled $a_i:a_j$ for $i, j \in \{1, 2, \dots, n\}$.

- The left subtree shows subsequent comparisons if $a_i < a_j$.
- The right subtree shows subsequent comparisons if $a_i \geq a_j$.

Decision-tree for insertion sort

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

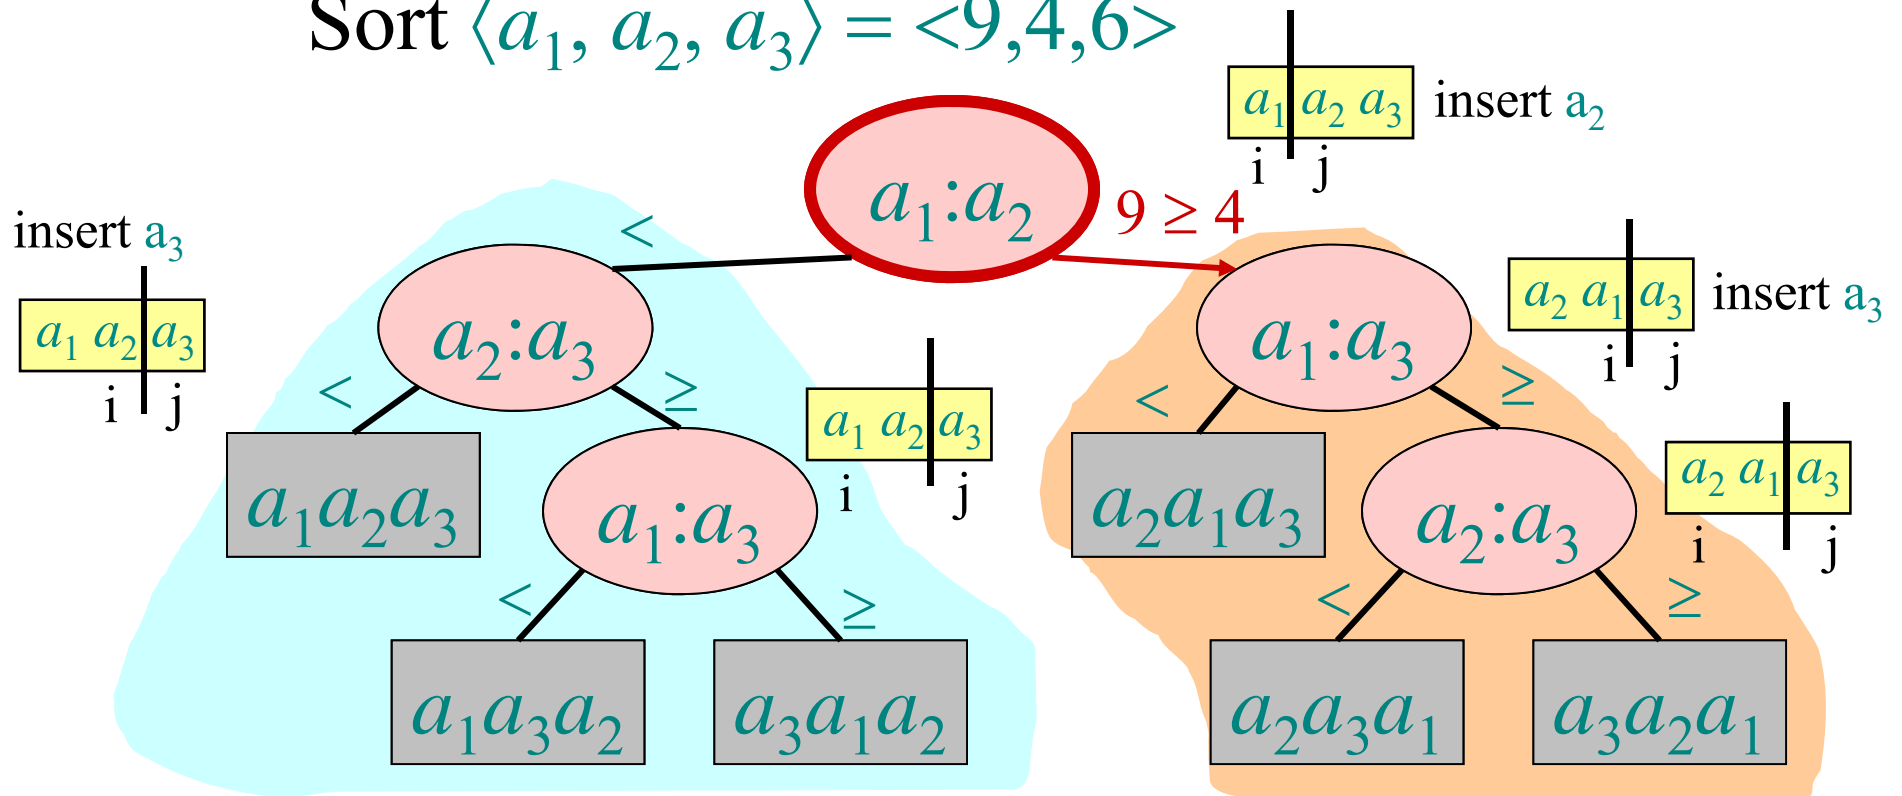


Each internal node is labeled $a_i:a_j$ for $i, j \in \{1, 2, \dots, n\}$.

- The left subtree shows subsequent comparisons if $a_i < a_j$.
- The right subtree shows subsequent comparisons if $a_i \geq a_j$.

Decision-tree for insertion sort

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

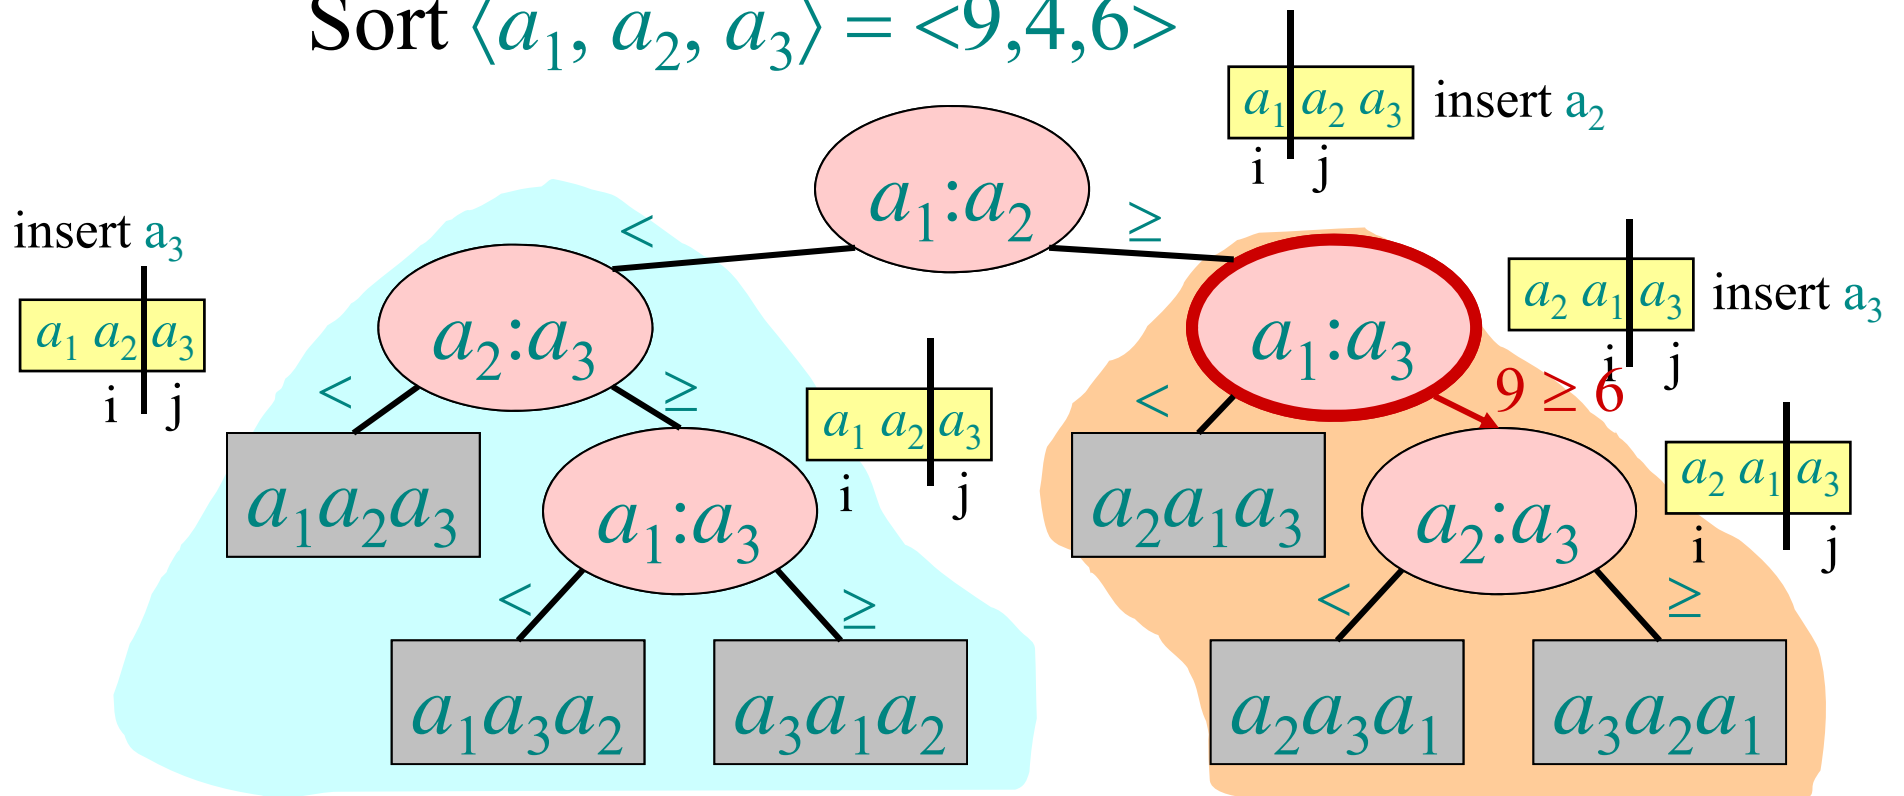


Each internal node is labeled $a_i:a_j$ for $i, j \in \{1, 2, \dots, n\}$.

- The left subtree shows subsequent comparisons if $a_i < a_j$.
- The right subtree shows subsequent comparisons if $a_i \geq a_j$.

Decision-tree for insertion sort

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

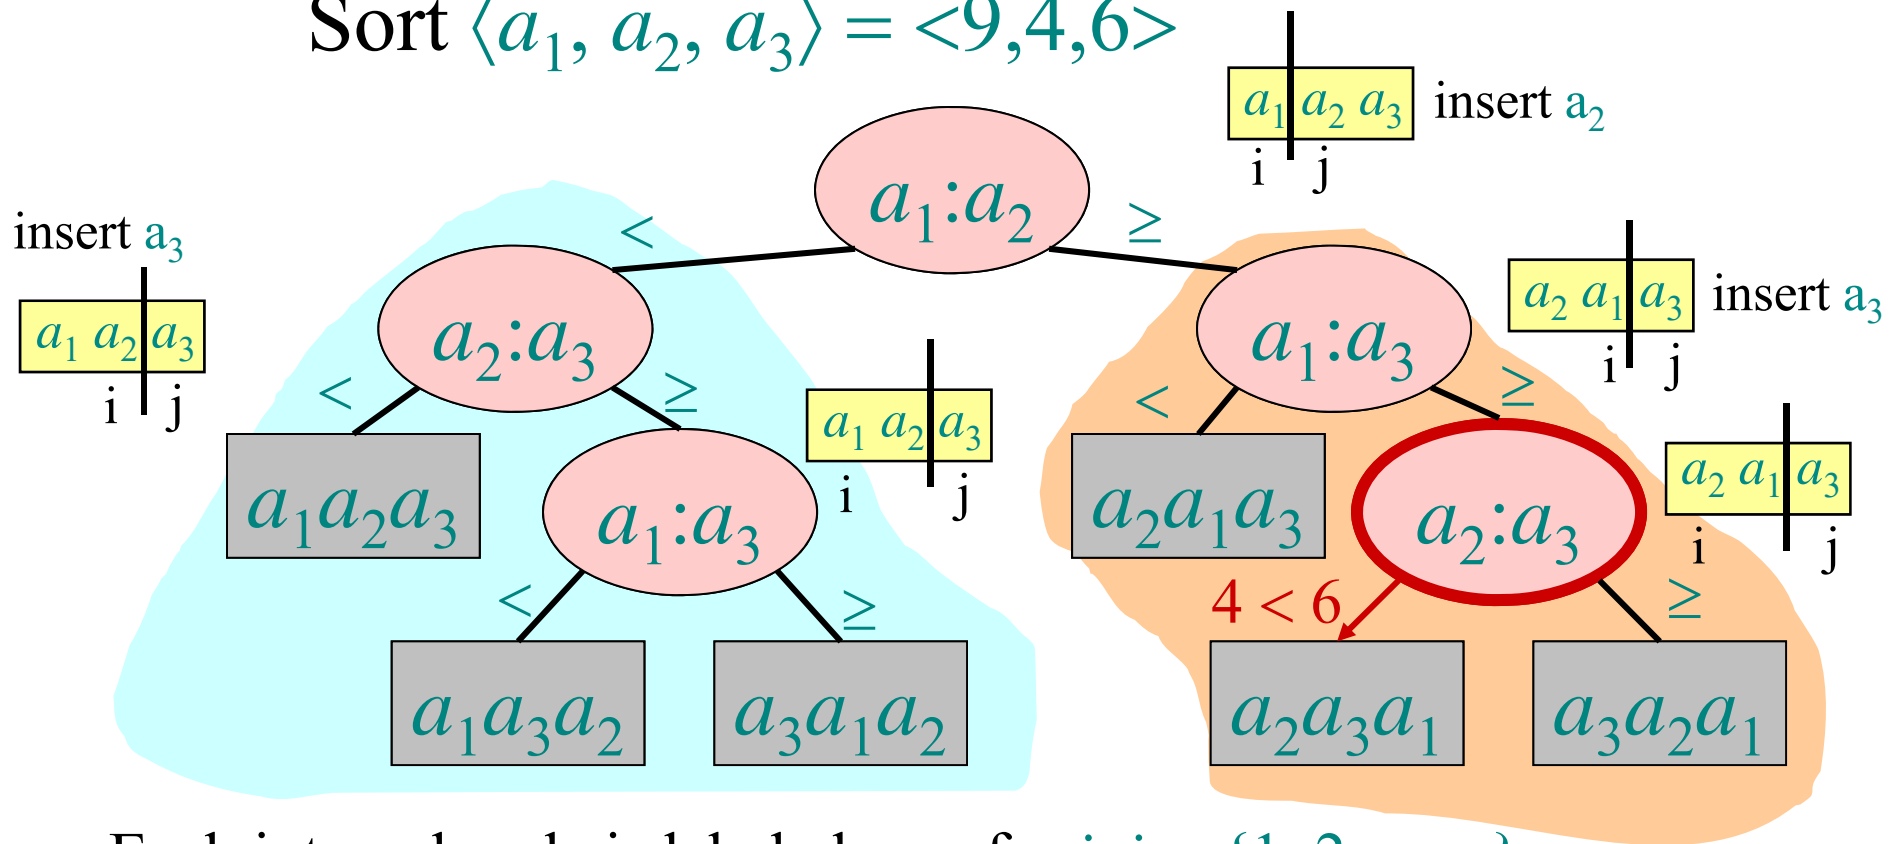


Each internal node is labeled $a_i:a_j$ for $i, j \in \{1, 2, \dots, n\}$.

- The left subtree shows subsequent comparisons if $a_i < a_j$.
- The right subtree shows subsequent comparisons if $a_i \geq a_j$.

Decision-tree for insertion sort

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

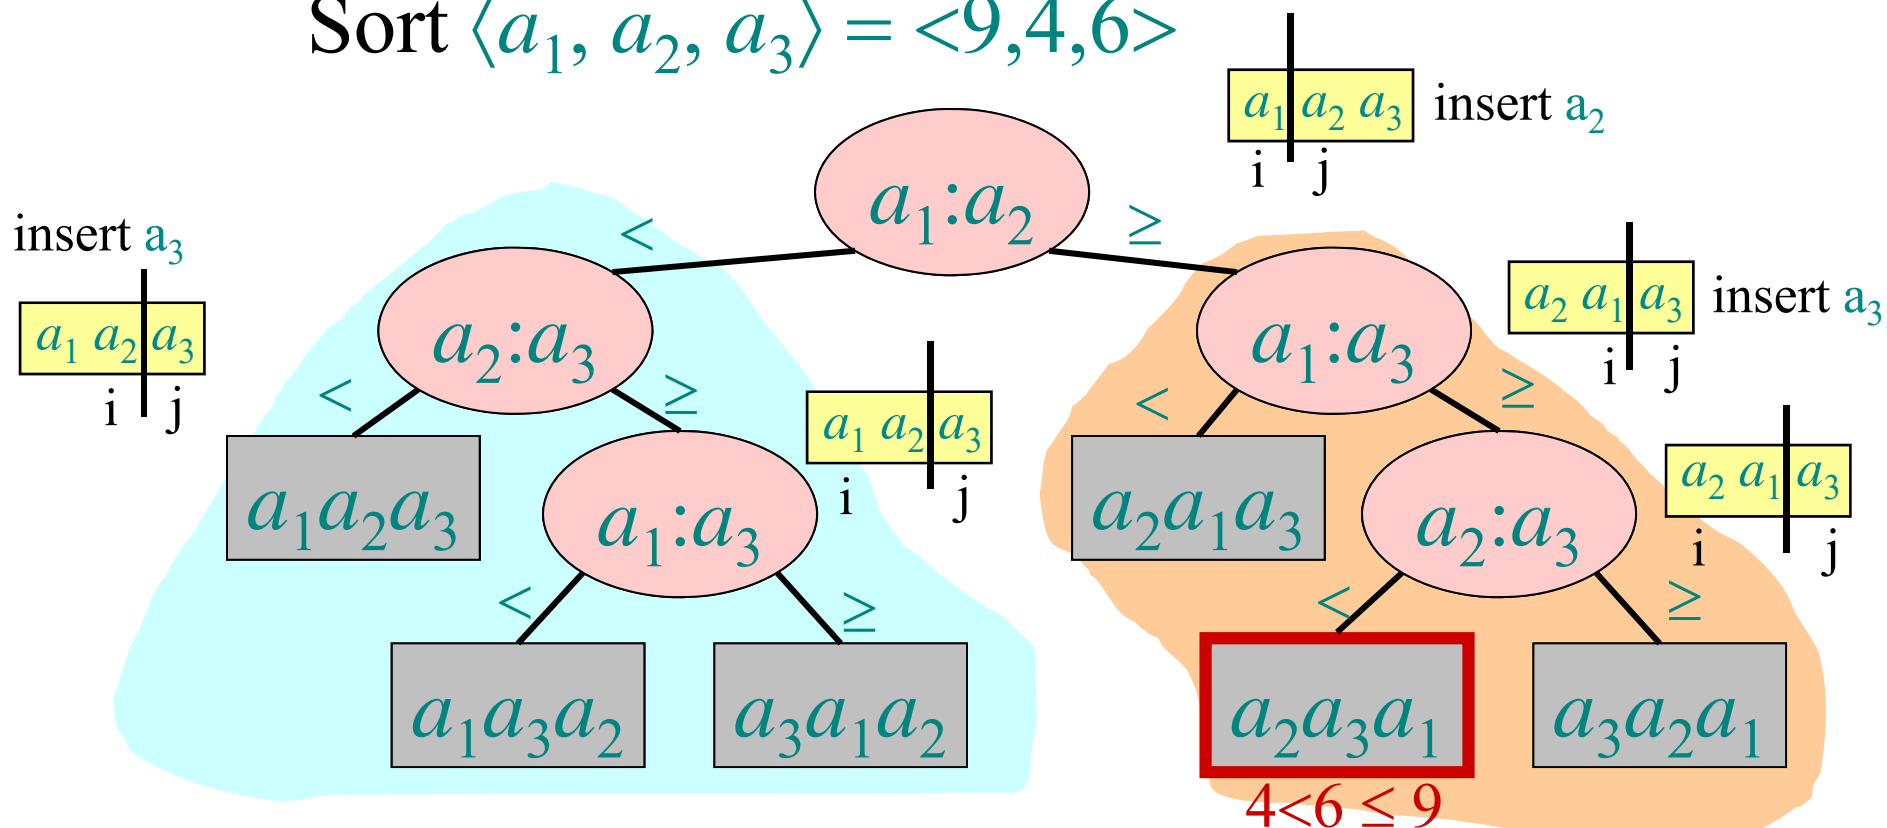


Each internal node is labeled $a_i:a_j$ for $i, j \in \{1, 2, \dots, n\}$.

- The left subtree shows subsequent comparisons if $a_i < a_j$.
- The right subtree shows subsequent comparisons if $a_i \geq a_j$.

Decision-tree for insertion sort

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

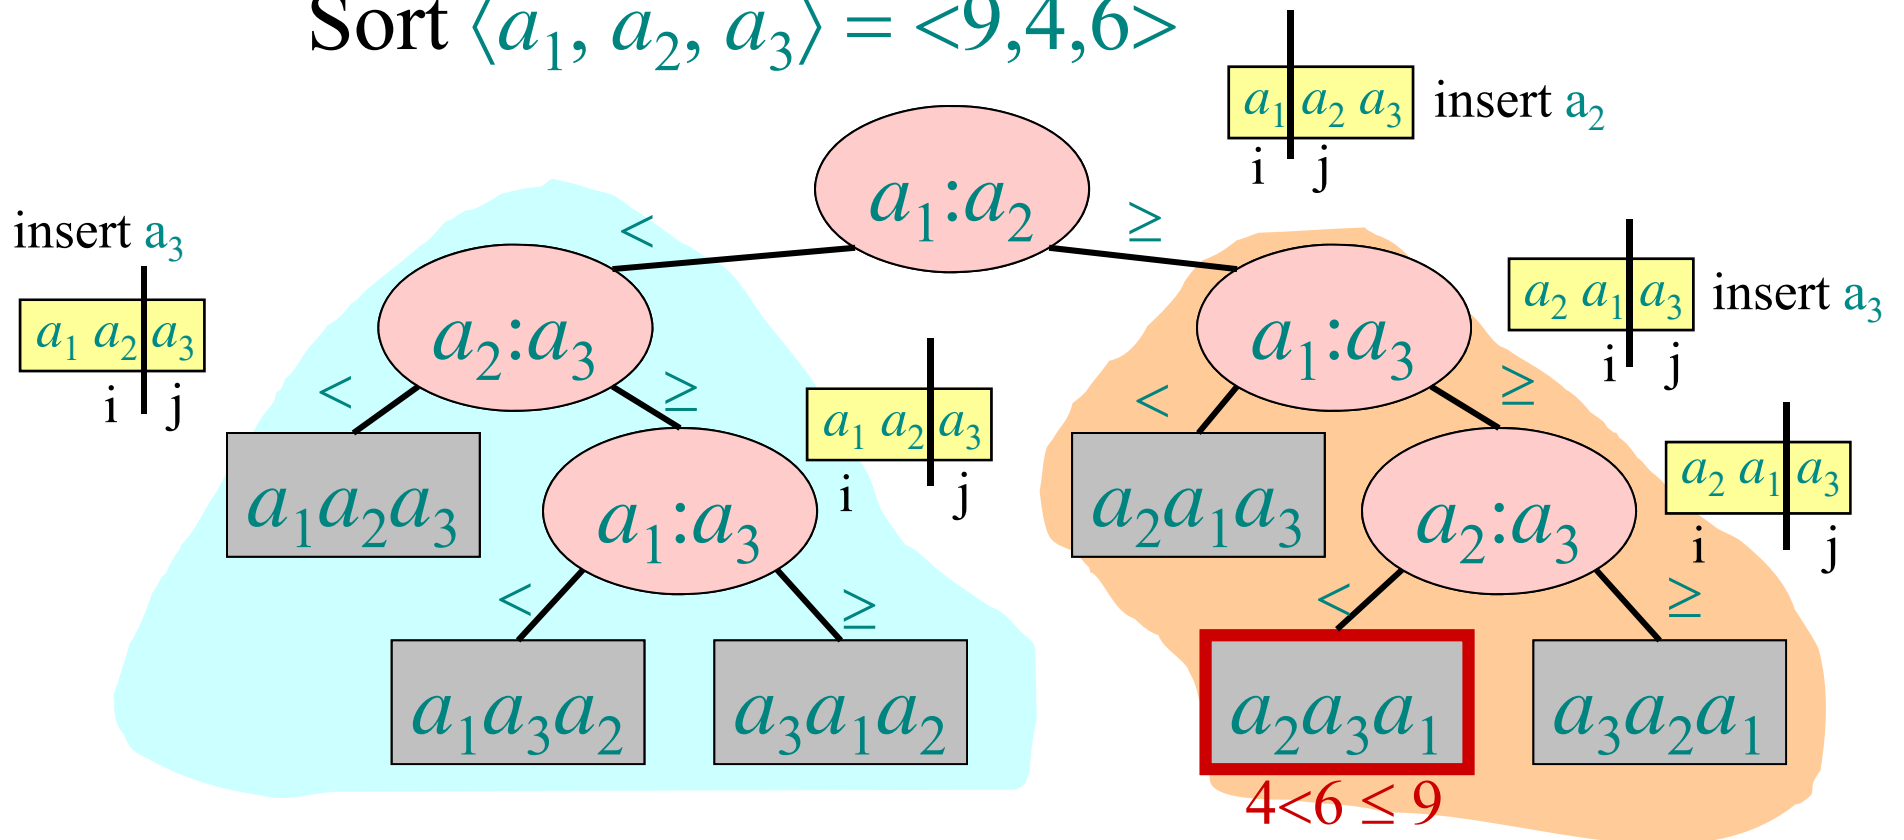


Each internal node is labeled $a_i:a_j$ for $i, j \in \{1, 2, \dots, n\}$.

- The left subtree shows subsequent comparisons if $a_i < a_j$.
- The right subtree shows subsequent comparisons if $a_i \geq a_j$.

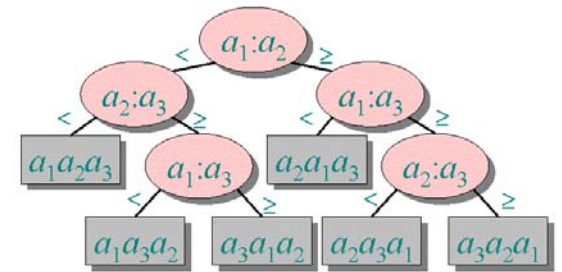
Decision-tree for insertion sort

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



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

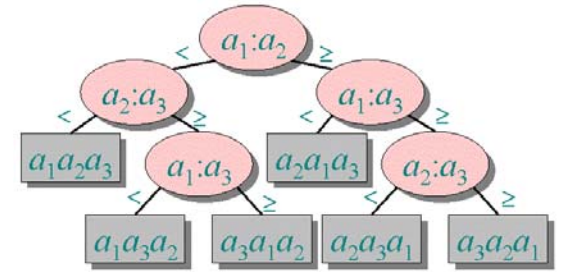
Decision-tree model



A decision tree models the execution of any comparison sorting algorithm:

- One tree per input size n .
- The tree contains **all** possible comparisons (= if-branches) that could be executed for **any** input of size n .
- The tree contains **all** comparisons along **all** possible instruction traces (= control flows) for **all** inputs of size n .
- For one input, only one path to a leaf is executed.
- Running time = length of the path taken.
- Worst-case running time = height of tree.

Lower bound for comparison 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. For a binary tree of height- h holds that $\#leaves \leq 2^h$. Thus, $n! \leq 2^h$.

$$\begin{aligned} \therefore h &\geq \log(n!) && (\log \text{ is mono. increasing}) \\ &\geq \log((n/2)^{n/2}) \\ &= n/2 \log n/2 \\ &\Rightarrow h \in \Omega(n \log n). \end{aligned}$$



Lower bound for comparison sorting

Corollary. Mergesort is an asymptotically optimal comparison sorting algorithm.



Sorting in linear time

Counting sort: No comparisons between elements.

- **Input:** $A[0 \dots n-1]$, where $A[j] \in \{0, 1, 2, \dots, k-1\}$.
- **Output:** $B[0 \dots n-1]$, sorted.
- **Auxiliary storage:** $C[0 \dots k-1]$.

Counting sort

1.for ($i = 0; i < k; i++$)

$C[i] = 0$

2.for ($j = 0; j < n; j++$)

$C[A[j]] = C[A[j]] + 1$ // $C[i] == |\{\text{key} = i\}|$

3.for ($i = 1; i < k; i++$)


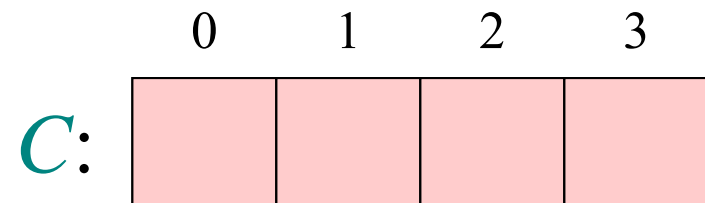
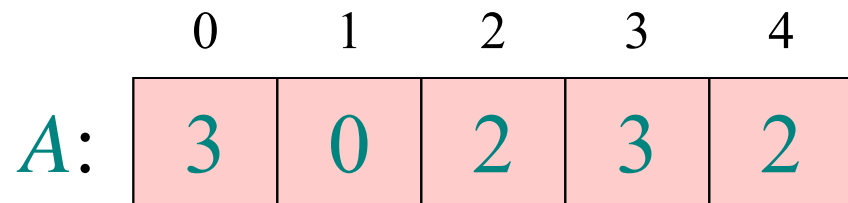
$C[i] = C[i] + C[i-1]$ // $C[i] == |\{\text{key} \leq i\}|$

4.for ($j = n-1; j \geq 0; j--$)

$B[C[A[j]]-1] = A[j]$

$C[A[j]] = C[A[j]] - 1$

Counting-sort example



Loop 1

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	0	0	0	0

<i>B</i> :					
------------	--	--	--	--	--

1. for ($i = 0; i < k; i++$)
 $C[i] = 0$

Loop 2

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	0	0	0	1


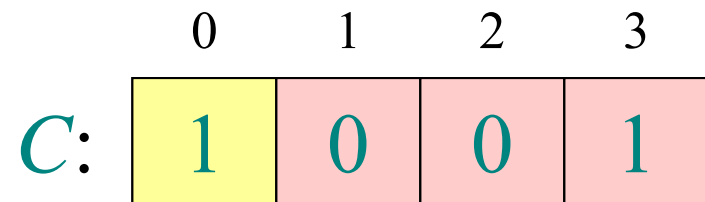
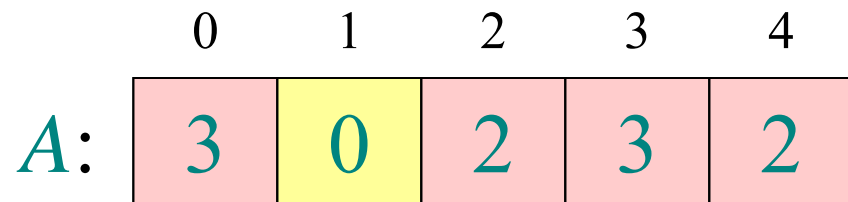
<i>B</i> :					
------------	--	--	--	--	--

2. **for** ($j = 0; i < n; j++$)

$C[A[j]] = C[A[j]] + 1$

// $C[i] == |\{\text{key} = i\}|$

Loop 2



2. for ($j = 0; i < n; j++$)

$C[A[j]] = C[A[j]] + 1$

// $C[i] == |\{\text{key} = i\}|$

Loop 2

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	1	0	1	1

<i>B</i> :					
------------	--	--	--	--	--

2. for ($j = 0; i < n; j++$)

$C[A[j]] = C[A[j]] + 1$

// $C[i] == |\{\text{key} = i\}|$

Loop 2

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	1	0	1	2

<i>B</i> :					
------------	--	--	--	--	--

2. **for** ($j = 0; i < n; j++$)

$C[A[j]] = C[A[j]] + 1$

$// C[i] == |\{\text{key} = i\}|$

Loop 2

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	1	0	2	2

<i>B</i> :					
------------	--	--	--	--	--

2. for ($j = 0; i < n; j++$)

$C[A[j]] = C[A[j]] + 1$

// $C[i] == |\{\text{key} = i\}|$

Loop 3

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	1	0	2	2

<i>B</i> :					
------------	--	--	--	--	--

<i>C'</i> :	1	1	2	2
-------------	---	---	---	---

3. for ($i = 1; i < k; i++$)

$$C[i] = C[i] + C[i-1]$$

$$// C[i] == |\{\text{key} \leq i\}|$$

Loop 3

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	1	0	2	2

<i>B</i> :					
------------	--	--	--	--	--

<i>C'</i> :	1	1	3	2
-------------	---	---	---	---

3. for ($i = 1; i < k; i++$)

$$C[i] = C[i] + C[i-1]$$

$$// C[i] == |\{\text{key} \leq i\}|$$

Loop 3

	0	1	2	3	4
<i>A</i> :	3	0	2	3	2

	0	1	2	3
<i>C</i> :	1	0	2	2

<i>B</i> :					
------------	--	--	--	--	--

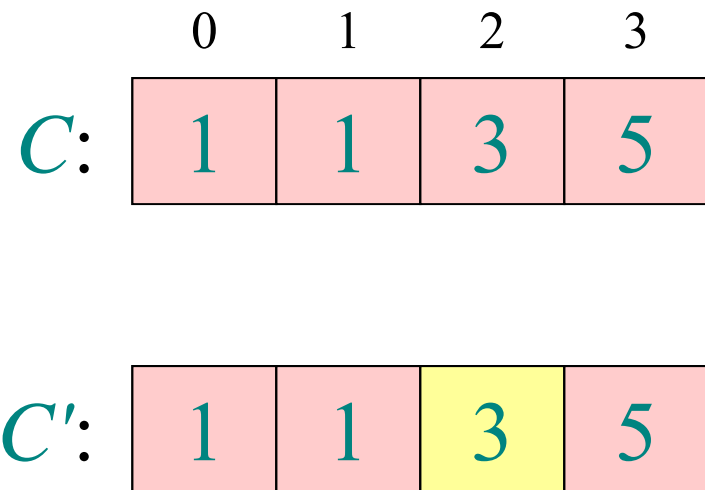
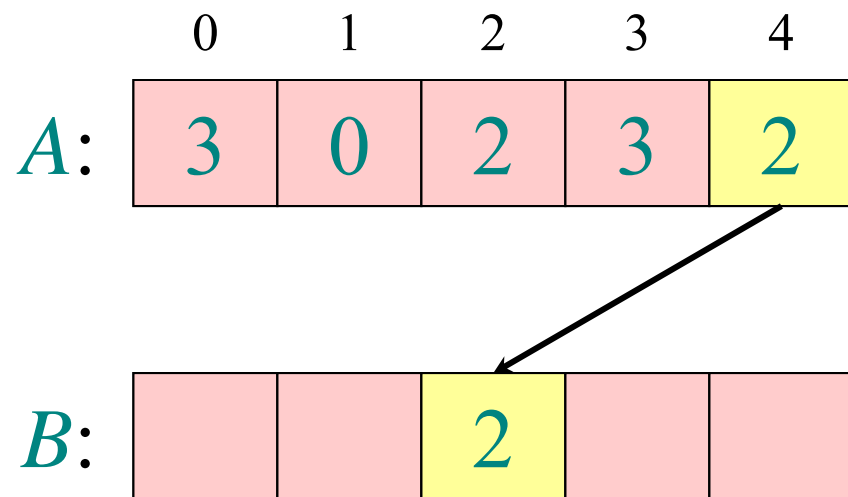
<i>C'</i> :	1	1	3	5
-------------	---	---	---	---

3. **for** ($i = 1; i < k; i++$)

$$C[i] = C[i] + C[i-1]$$

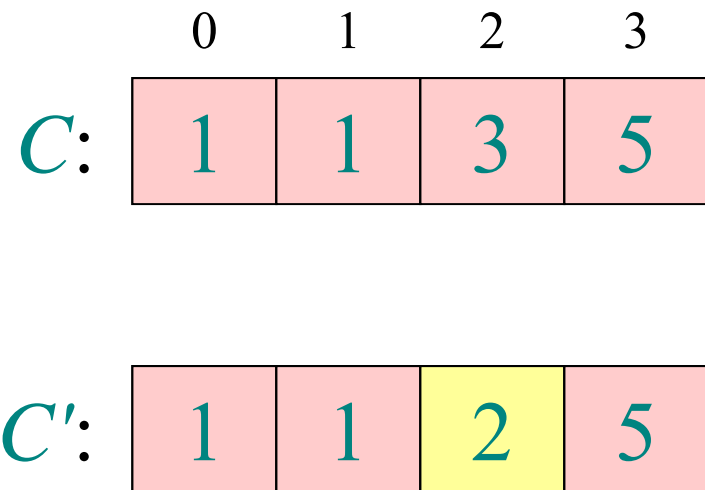
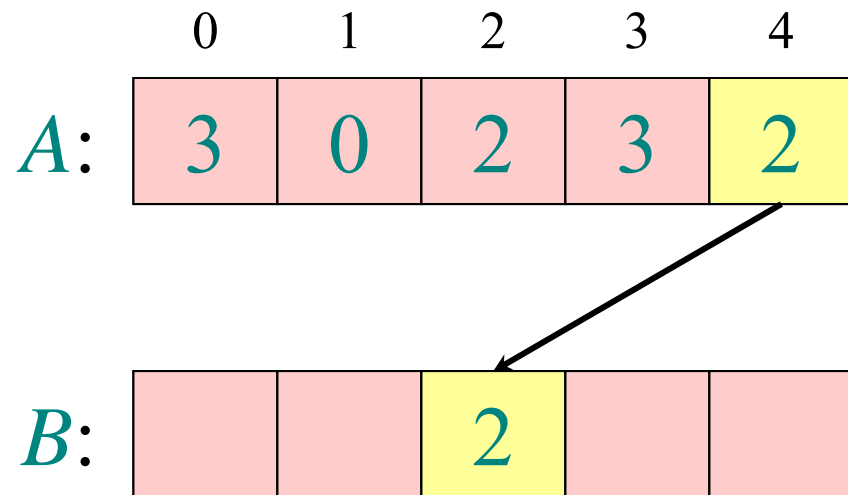
$$// C[i] == |\{\text{key} \leq i\}|$$

Loop 4



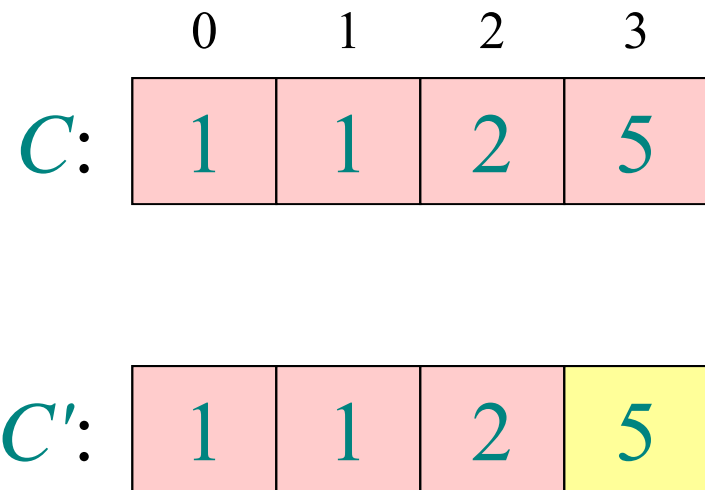
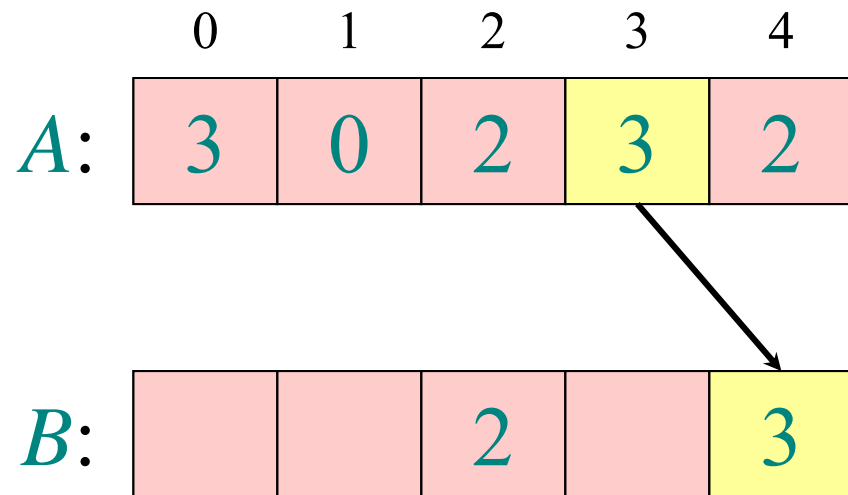
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



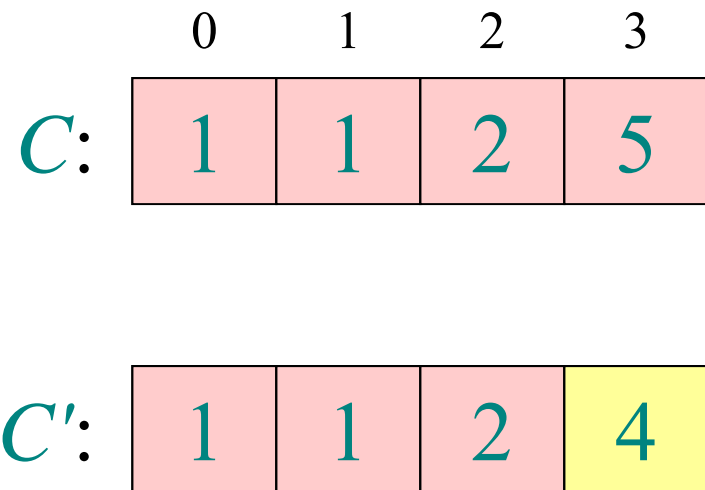
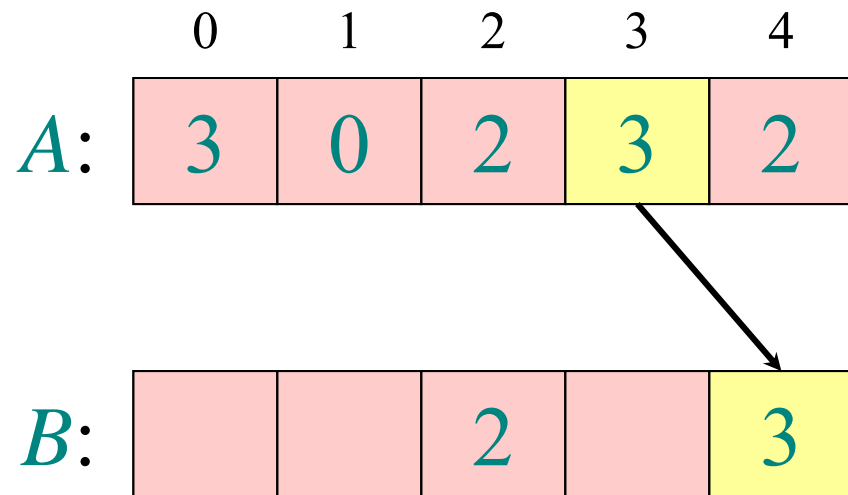
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



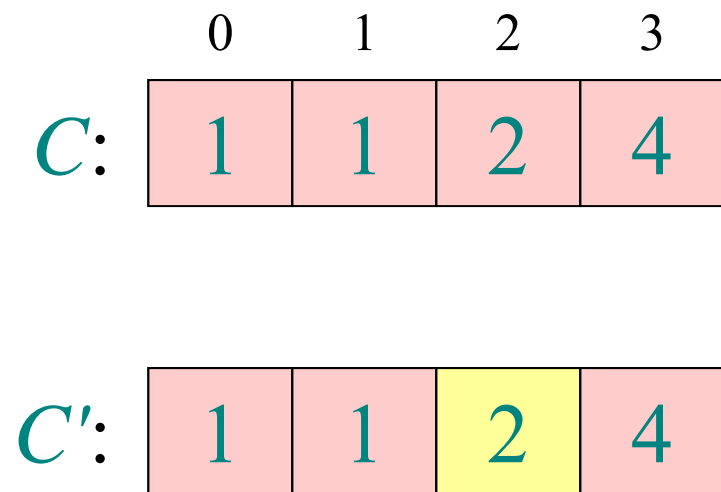
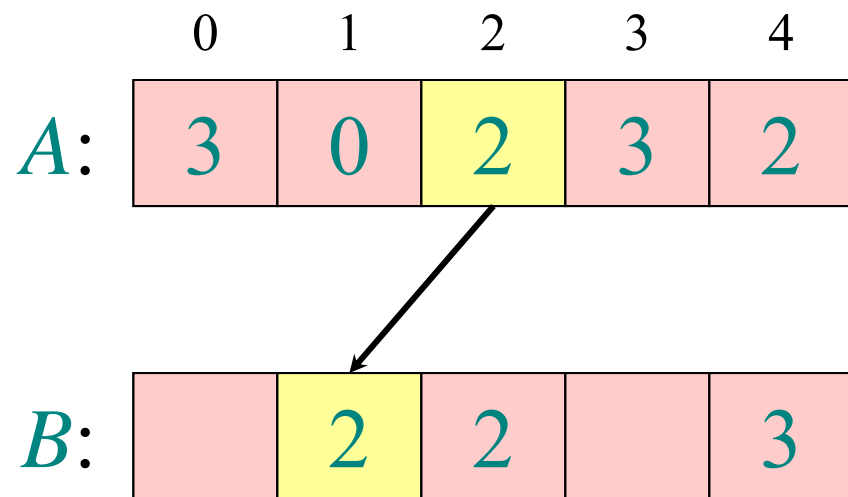
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



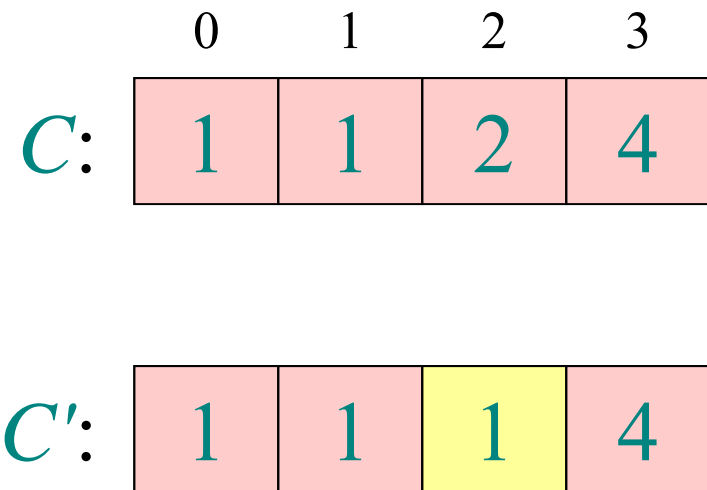
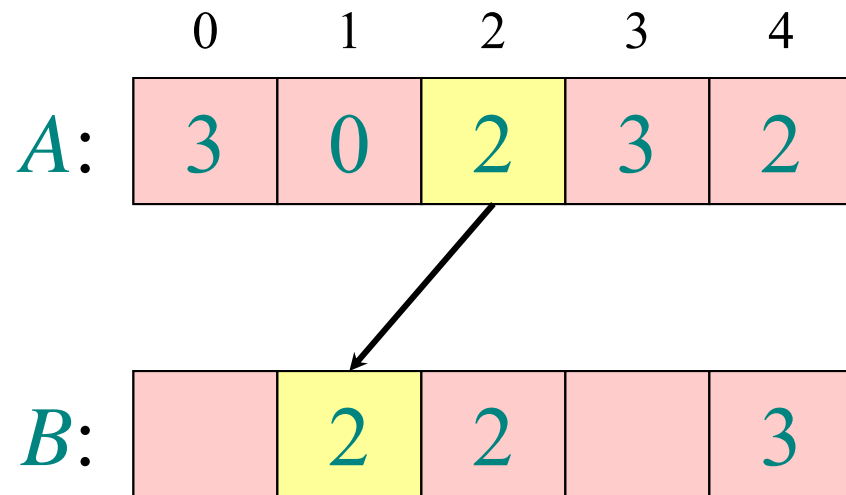
4. **for** ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



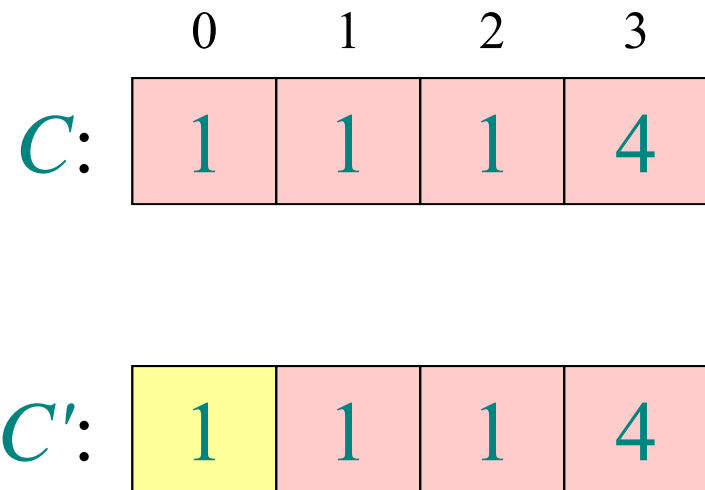
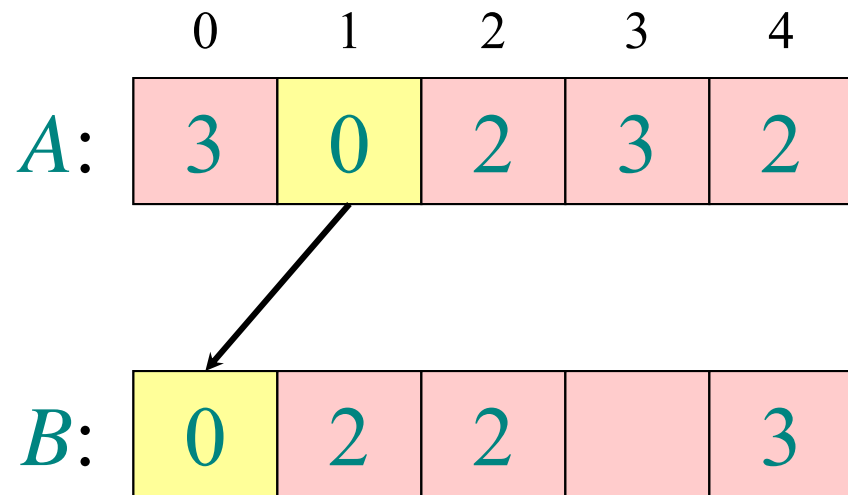
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



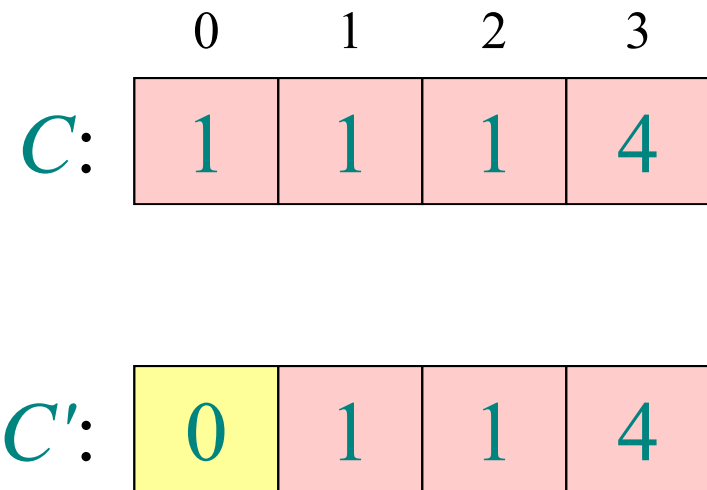
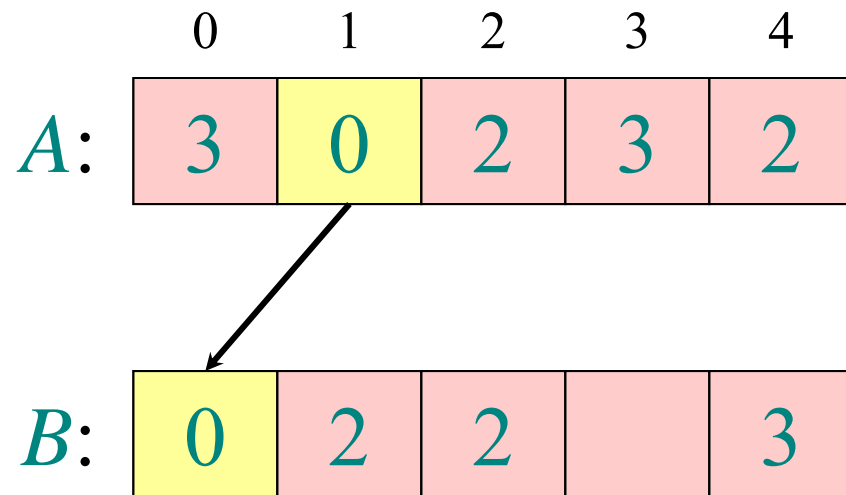
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



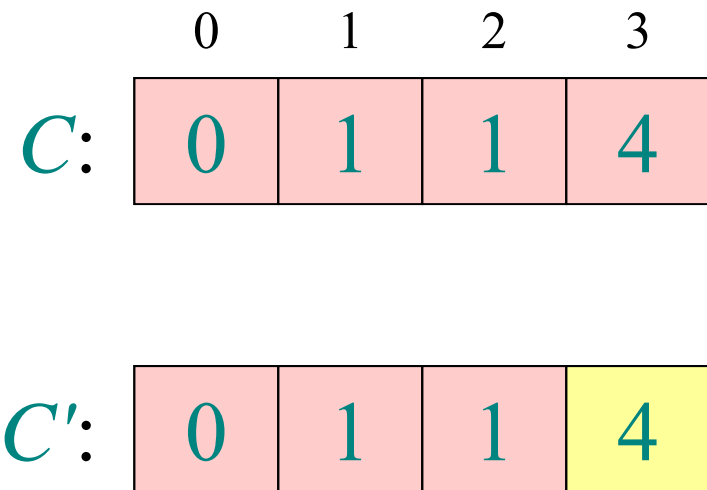
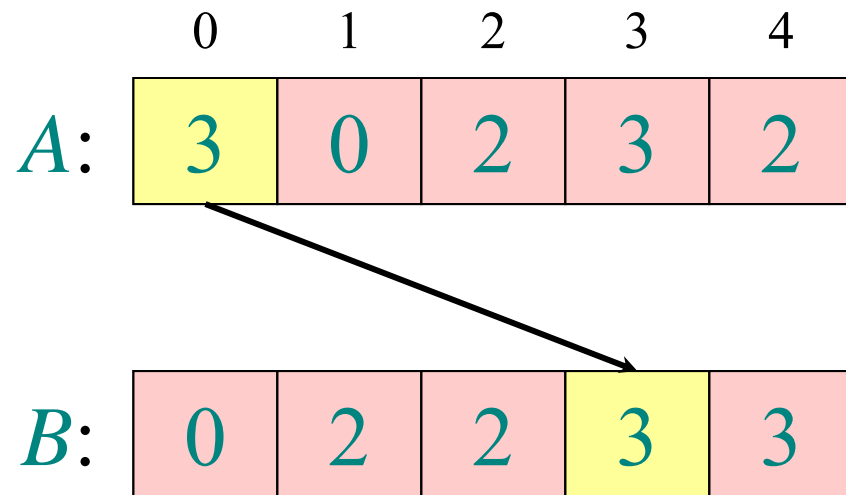
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



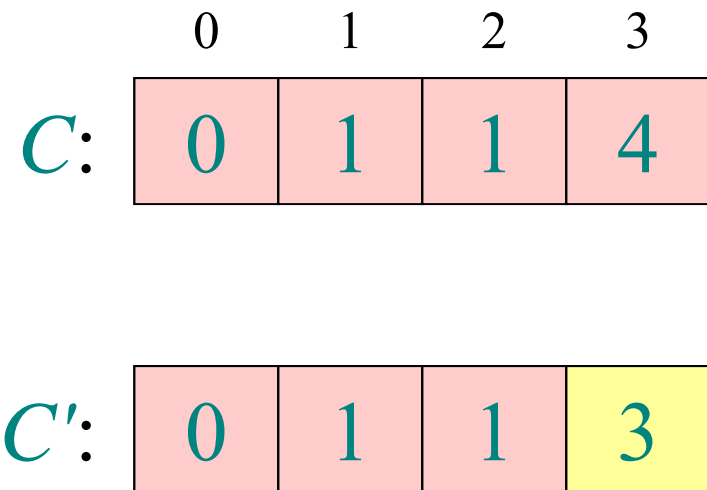
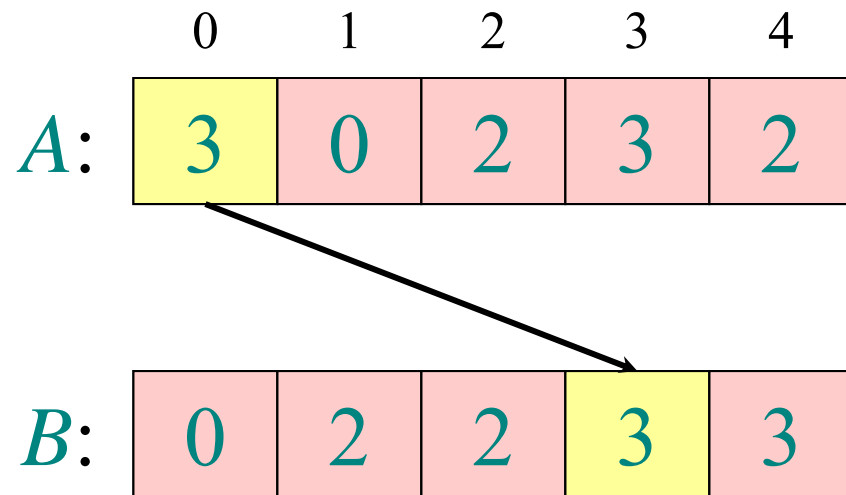
4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Loop 4



4. for ($j = n-1; i \geq 0; j--$)
 $B[C[A[j]]-1] = A[j]$
 $C[A[j]] = C[A[j]] - 1$

Analysis

$$\Theta(k) \left\{ \begin{array}{l} \mathbf{1. for} (i = 0; i < k; i++) \\ C[i] = 0 \end{array} \right.$$

$$\Theta(n) \left\{ \begin{array}{l} \mathbf{2. for} (j = 0; j < n; j++) \\ C[A[j]] = C[A[j]] + 1 \end{array} \right.$$

$$\Theta(k) \left\{ \begin{array}{l} \mathbf{3. for} (i = 1; i < k; i++) \\ C[i] = C[i] + C[i-1] \end{array} \right.$$

$$\Theta(n) \left\{ \begin{array}{l} \mathbf{4. for} (j = n-1; j \geq 0; j--) \\ B[C[A[j]]-1] = A[j] \\ C[A[j]] = C[A[j]] - 1 \end{array} \right.$$

$$\Theta(n + k)$$

Running time

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

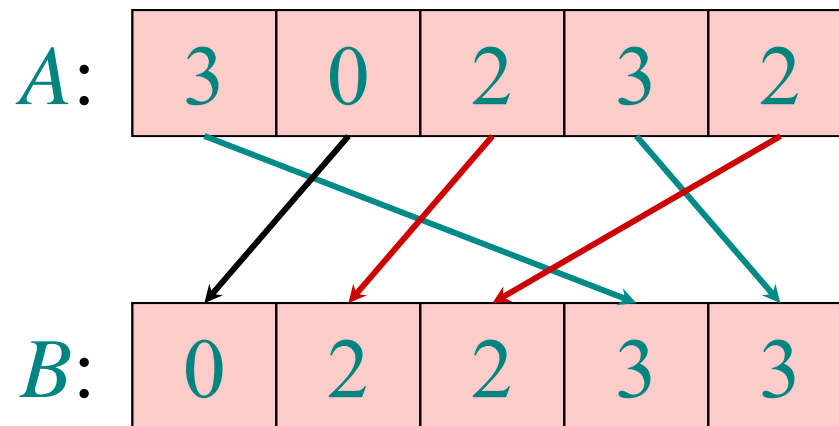
- But, sorting takes $\Omega(n \log n)$ time!
- Where's the fallacy?

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.



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 (left to right).
- Good idea: Sort on *least-significant digit first* (right to left) with an auxiliary *stable* sorting algorithm (like counting sort).

3 2 9

4 5 7

6 5 7

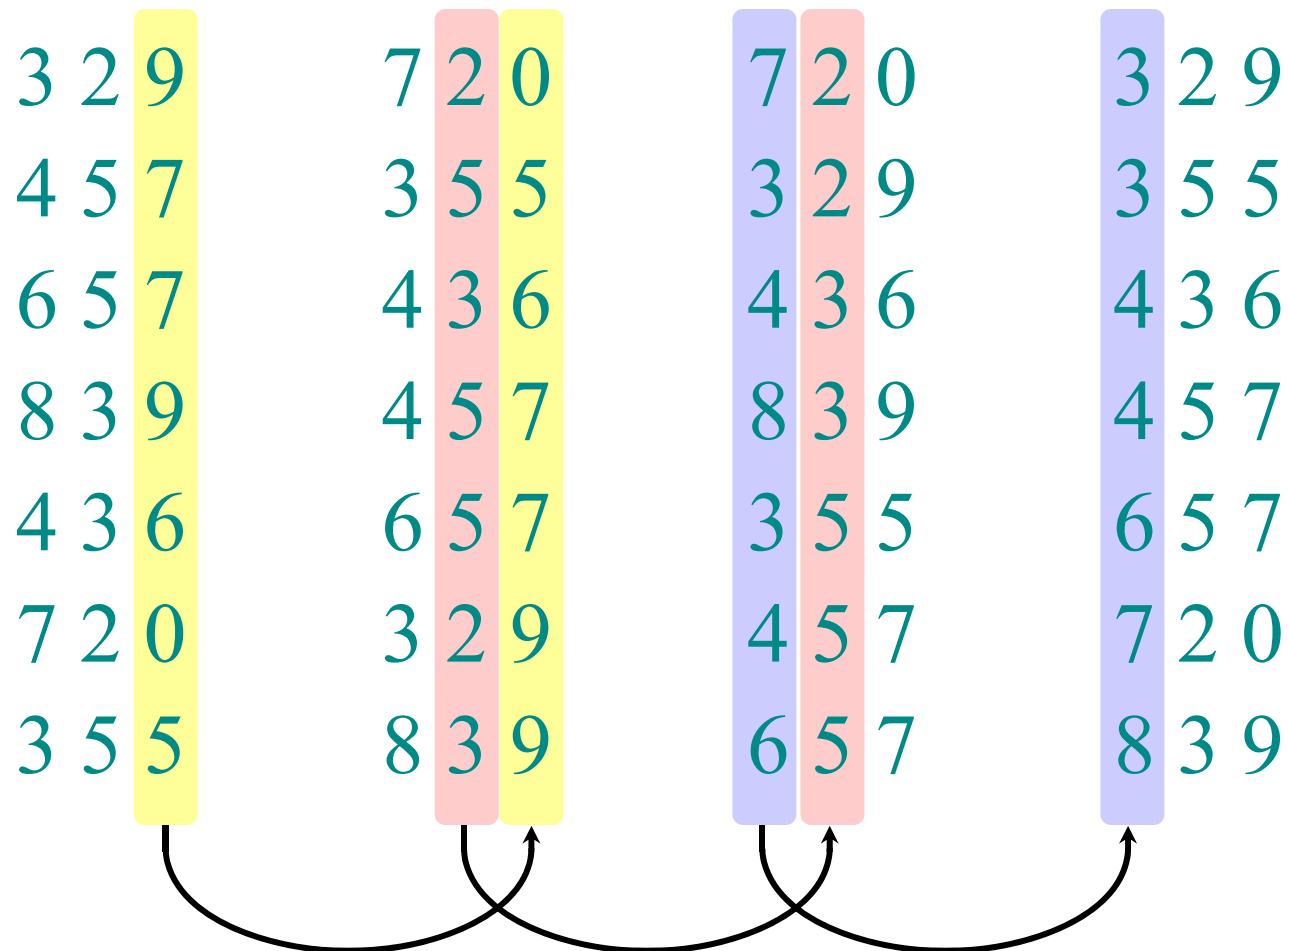
8 3 9

4 3 6

7 2 0

3 5 5

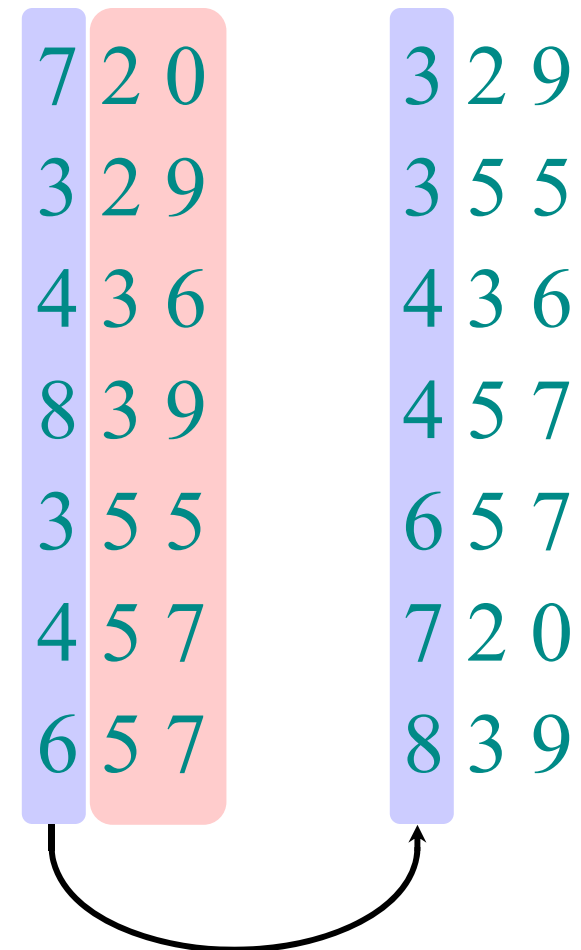
Operation of radix sort



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



Correctness of radix sort

Induction on digit position

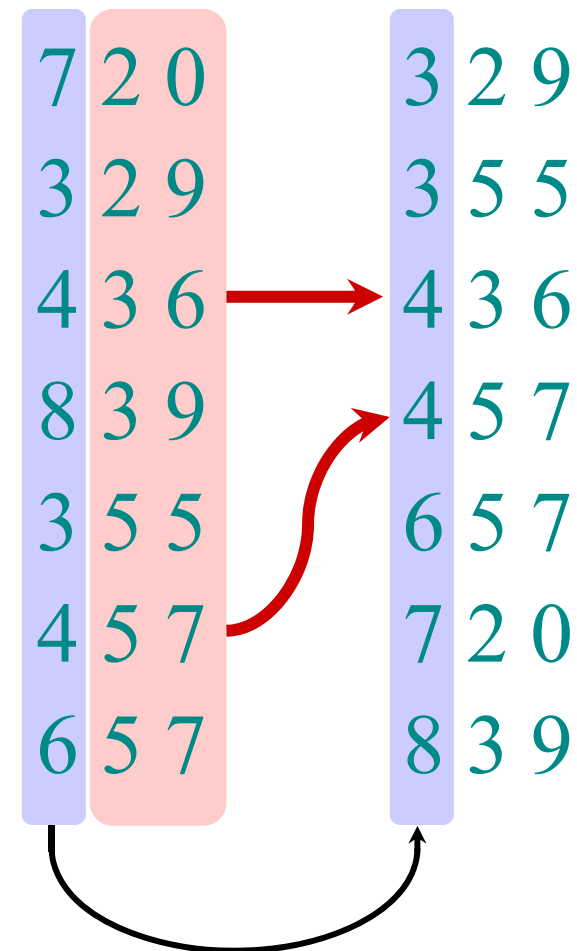
- Assume that the numbers are sorted by their low-order $t - 1$ digits.
- Sort on digit t
 - Two numbers that differ in digit t are correctly sorted.



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.

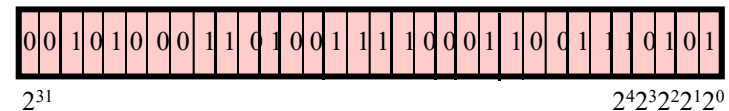


Analysis of radix sort

- Sort n computer words of b bits each.
- View each word as having b/r base- 2^r digits.

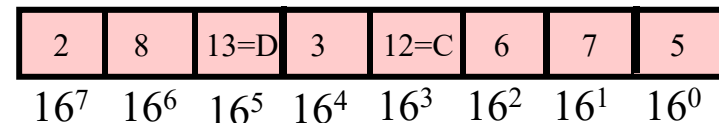
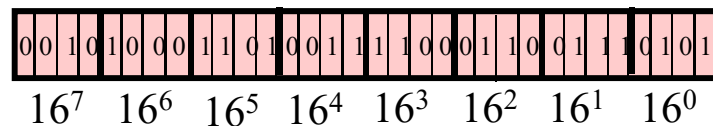
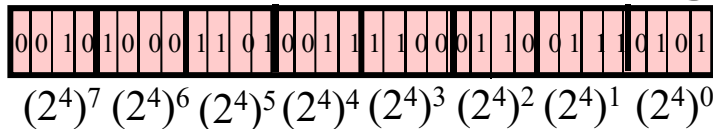
Example: 32-bit word ($b=32$)

- $r = 1$: 32 base-2 digits



$\Rightarrow b/r = 32$ passes of counting sort on base-2 digits

- $r = 4$: 32/4 base- 2^4 digits (hexadecimal numbers)

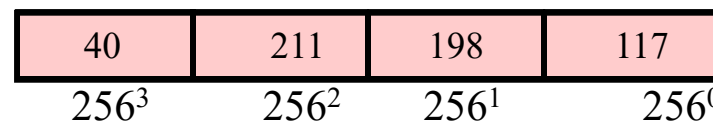
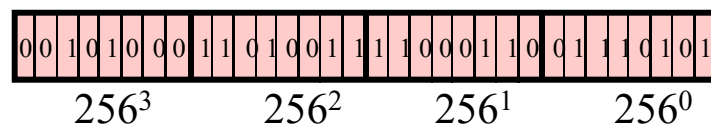
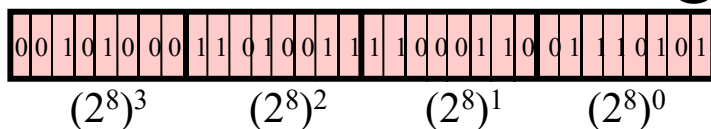


$\Rightarrow b/r = 8$ passes of counting sort on base- 2^4 digits

Analysis of radix sort (cont.)

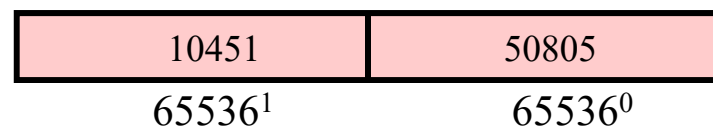
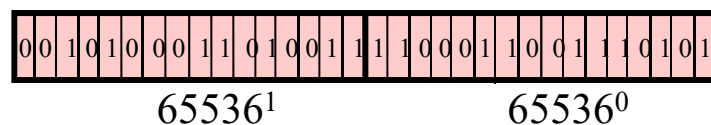
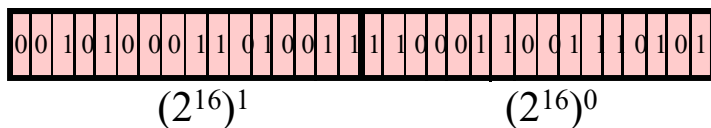
Example: 32-bit word ($b=32$)

- $r = 8$: $32/8$ base- 2^8 digits



$\Rightarrow b/r = 4$ passes of counting sort on base- 2^8 digits

- $r = 16$: $32/16$ base- 2^{16} digits



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

Analysis of radix sort

- Sort n computer words of b bits each.
- View each word as having b/r base- 2^r digits.
- Assume counting sort is the auxiliary stable sort.
- Make b/r passes of counting sort on base- 2^r 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

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

- Choose r to minimize $T(n, b)$:
Increasing r means fewer passes, but as $r \gg \log n$, the time grows exponentially.

Choosing r

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

Minimize $T(n, b)$:

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 $T(n, b) = \Theta(bn/\log n)$.

Radix Sort with optimized r

- 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 $\Theta(dn)$ 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):

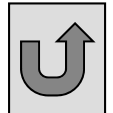
- At most 3 passes when sorting ≥ 2000 numbers.
- Merge sort and quicksort do at least $\lceil \log 2000 \rceil = 11$ passes.

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.



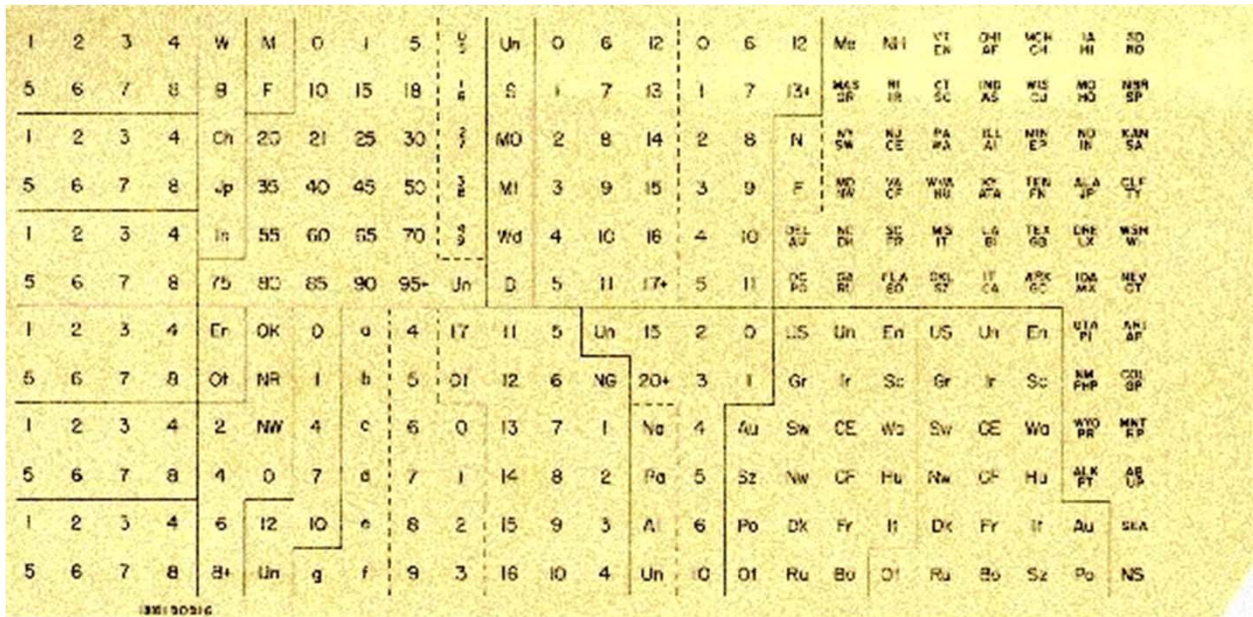
Herman Hollerith (1860-1929)



- 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

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

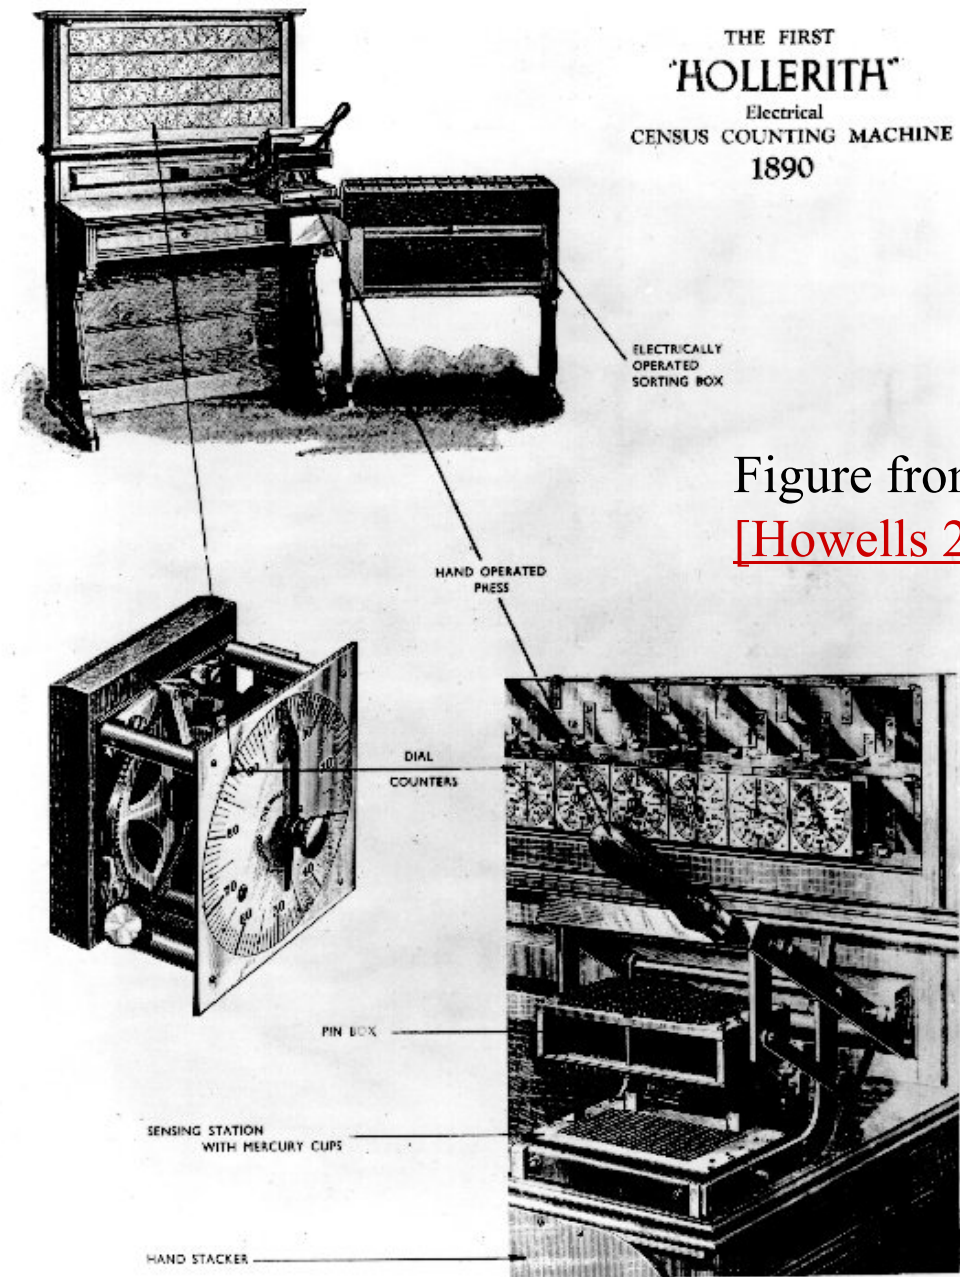



Figure from [\[Howells 2000\]](#).

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.
- 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.



Hollerith Tabulator, Pantograph, Press, and Sorter

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!