

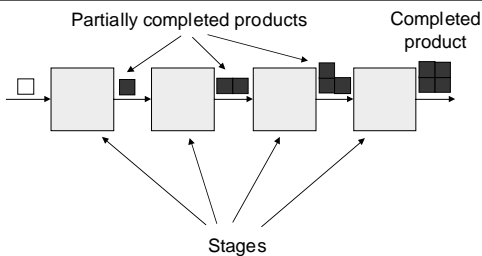
Pipelining

PP Ch. 5

Pipelining Defined


- A “product” in a stage of partial completion, is moved along through a series of “stations”, becoming more complete at each station.
- The stations operate in parallel on multiple products, each working on a product in its respective stage of completion.
- There may or may not be parallelism *within* a given station.

Pipelining



Examples of Computational Products


Product	Partially Complete Product
Number	Approximation to a number
Sorted array	Partially sorted array
Image	Transformed image




Utility of Pipelines

- Pipelines can provide a modular approach to constructing functions, rather than trying to invent or manage one single comprehensive function.
- Unix users are used to this kind of thing:
`(thin < image) | rotate | crop | invert`

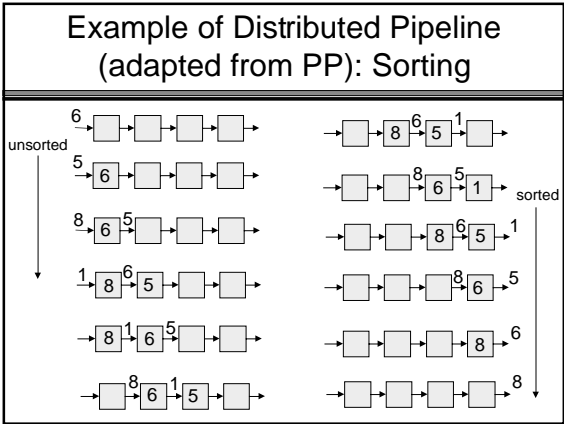
Pipeline Timing from the products' view





Distributed Pipelining (my definition)

- In *distributed* pipelining, the result is not what comes out the last stage in one step, but rather the **collective sequence** of things coming out.

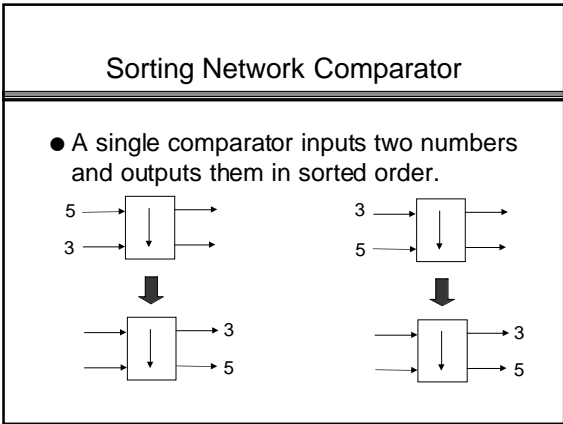


Pipeline Sorting

- The method shown is $O(N)$ and requires N processors.
- Also, the number of items to be sorted is limited by the number of processors.
- As soon as a sequence clears the first processor, the next sequence can be started, so $N-1$ processors can be kept busy.


Pipelined, Parallel, Sorting Methods

- Sorting networks were first proposed by Kenneth Batcher (Goodyear Aerospace) in the 1960's.
- They are networks constructed from a number of simple devices, called **comparators**.
- Sorting networks are discussed in PP 9.2.7 and 9.2.8. We discuss them here with pipelining.




Sorting Network Comparator

- If desired, an individual comparator *can* be pipelined at the *bit-level*, as a finite-state machine accepting inputs MSB first (assuming the same number of bits in both numbers):




Sorting Networks


- Comparator abstraction (for drawing networks)



- A network that sorts 3 numbers



Analysis




Exercises

- Construct a sorting network for 4 numbers
- Construct a sorting network for 8 numbers

General Constructions

- Odd-Even Merging
 - Assume $2n$ elements to be sorted.
 - Split the elements into two groups.
 - Sort each half recursively.
 - Merge the odd-indexed components of the result.
 - Merge the even-indexed components of the result.
 - Merge the output of the merges.
- parallel {
- can be done in 1 stage!!

Merging



In general, how can one verify that a proposed scheme sorts?

- 0-1 Principle
- Symbolic analysis (“Boolean” operators)
- Partial-order analysis (Gale & Karp, Liu)

Timing Analysis of Sorting by Odd-Even Merging

- Let $S(n)$ = parallel time to sort n elements
- Let $M(n)$ = time to merge $n/2$ with $n/2$ elements
- Recurrences:
 - $S(1) = 0$
 - $S(2n) = S(n) + M(2n)$
 - $M(2) = 1$
 - $M(2n) = M(n) + 1$
- assuming we don't charge for splitting into 2 groups


Solving Recurrences: Assume powers of 2

- Merging:
 - $M(2) = 1$
 - $M(2n) = M(n) + 1$
 - Therefore
 - $M(2^n) = n$
 - Sorting:
 - $S(2^n) = S(2^{n-1}) + n$
 - Giving $S(2^n) = n + (n-1) + (n-2) + \dots + 1 = n^*(n+1)/2$
- i.e. $S(N) = O(\log^2 N)$


A Different Sort of Construction

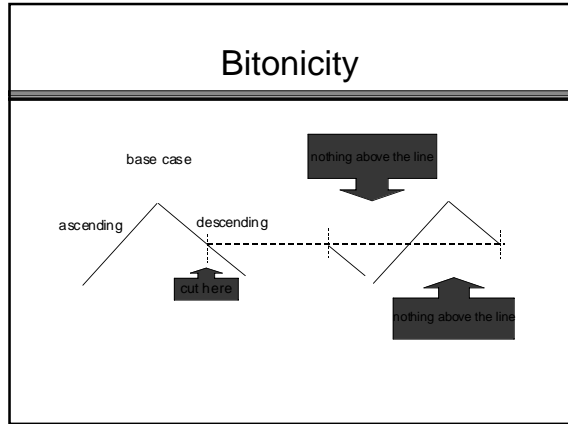
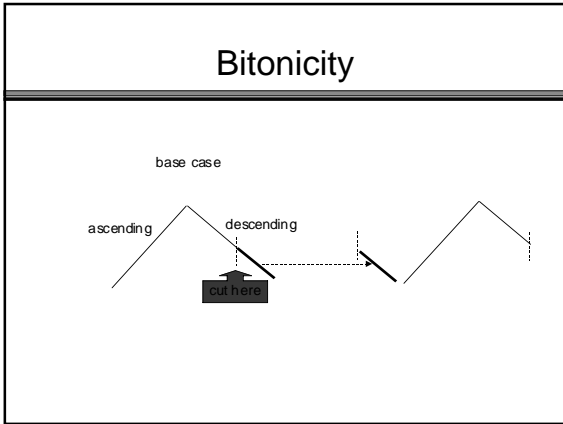
- Bitonic sorting
- Define a bitonic sequence to be one in which is either:
 - a strictly ascending portion, followed by a strictly descending portion, OR
 - any cyclic shift of a bitonic sequence.

Bitonicity



Bitonicity





Bitonic Sorting:

(This is the main *reason* for the definition of bitonic.)

- An already-bitonic sequences can be sorted by the following scheme:
 - Do pairwise compare-exchanges between elements j and $j+n/2$ for $j = 1, 2, \dots, n/2$
 - This gives two sequences that:
 - are bitonic
 - all elements of one are \leq all elements of the other
 - Repeating this scheme recursively with the resulting two bitonic sequences gives a sorted sequence

A Point that Needs Proof

- all elements of one are \leq all elements of the other
- In order to show this, we could use the 0-1 principle.
- Map an arbitrary numeric sequence into 0's and 1's by choosing an arbitrary threshold θ :
 $x > \theta$ maps x to 1, otherwise x is mapped to 0.
- For a bitonic sequence, the corresponding 0-1 sequence is 000...01...1 or **some rotation** thereof.
- When we do the compare-exchange step, the result will have the desired property.
- Example: $1111000011 \Rightarrow 0001011111$

A Point that Needs Proof, cont'd

- Example: $1111000011 \Rightarrow 0001011111$

If the number of 0's and 1's is the same, then we end up with the form:

$$0 \dots 0 \ 1 \dots 1$$

If there are more 0's, we end up with the form

$$0 \dots 0 \ 0 \ . \ 1 \ . \ 1 \ . \ 0$$

If there are more 1's, we end up with the form

$$0 \ . \ 1 \ . \ 1 \ . \ 0 \ 1 \ . \ . \ . \ 1$$

all of which have the desired property.

Bitonic Sorting Example

- Start with: [8, 9, 7, 4, 2, 1, 3, 5]
- Do compare-exchanges between elements j and $j+n/2$ for $j = 1, 2, \dots, n/2$:

This gives two sequences that:

- are bitonic
- all elements of one are \leq all elements of the other

- Repeating this scheme recursively with the resulting two bitonic sequences gives a sorted sequence. just like in quicksort.

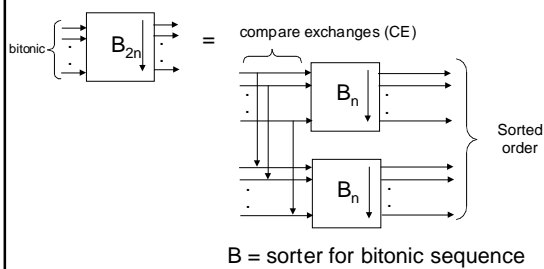
Bitonic Sorting

- Two already *sorted* sequences can be made into a bitonic sequence by reversing the second one and abutting them.
- This may seem like a counter-intuitive thing to do, since so arranging them is contrary to the desired ultimate sorted order.


Bitonic Sort Summary

- To create a bitonic sequence:
 - Sort the two halves of the sequence by any method, so that one half is ascending, the other descending.
- To sort a sequence known to be bitonic:
 - Compare-exchange elements j and $j+n/2$, for $j = 0, 1, 2, \dots, n/2-1$
 - It can be shown that the two halves are bitonic, and all of one half is \leq all of the other.
 - Therefore, sort them as bitonic sequences, then concatenate the results.


Bitonic Sort Recursion




General Sort Recursion



Bitonic Sort Examples



General Sort Examples



Proof of Shearsort

- Use the 0-1 principle
- Suffices to show that any initial assignment of 0's and 1's gets sorted.
- A row is called
 - 1-row, if it consists of all 1's
 - 0-row, if it consists of all 0's
 - mixed row, otherwise
- Need to show that sorting leaves us with 0-rows, a mixed row, then 1-rows.

Proof of Shearsort

- Need to show that sorting leaves us with 0-rows, a mixed row, then 1-rows.
- Claim: Each iteration at least halves the number of mixed rows.
- Thus after $\log n$ iterations ($n = \#rows$), there is only one mixed row.
- The final row sort sorts this mixed row, wherever it may be

Proof of Claim

- Claim: Each iteration at least halves the number of mixed rows.
- Proof: Consider an adjacent pair of mixed rows after a row sort.
- There can only be three kinds of configurations:
 - 0...1...1 equal 0's and 1's
1...10...0
 - 0...01...1 surplus of 0's
1...10...0
 - 0...01...1 surplus of 1's
1...10...0

Proof of Claim, cont'd

- Transitions from the three kinds of configurations:
 - $$\begin{array}{l} 0 \dots 1 \dots 1 \\ 1 \dots 10 \dots 0 \end{array} \rightarrow \begin{array}{l} 0 \dots \dots 0 \\ 1 \dots \dots 1 \end{array}$$
 - $$\begin{array}{l} 0 \dots 01 \dots 1 \\ 1 \dots 10 \dots 0 \end{array} \rightarrow \begin{array}{l} 0 \dots \dots 0 \\ 1 \dots 10 \dots 0 \end{array}$$
 - $$\begin{array}{l} 0 \dots 01 \dots 1 \\ 1 \dots \dots 10 \dots 0 \end{array} \rightarrow \begin{array}{l} 0 \dots 01 \dots 1 \\ 1 \dots \dots 1 \end{array}$$
- In each case, the number of mixed rows is at least halved, as desired.

Exercises

- Give an upper-bound for shared-memory implementation of shearsort.
- Give an upper-bound for distributed-memory implementation of shearsort, assuming a mesh-wise interconnection of PE's.
- Extend shearsort to 3 dimensions, and analyze.