

§9 Parallel Algorithms

*Design & Analysis
of Algorithms
Martin Ziegler*

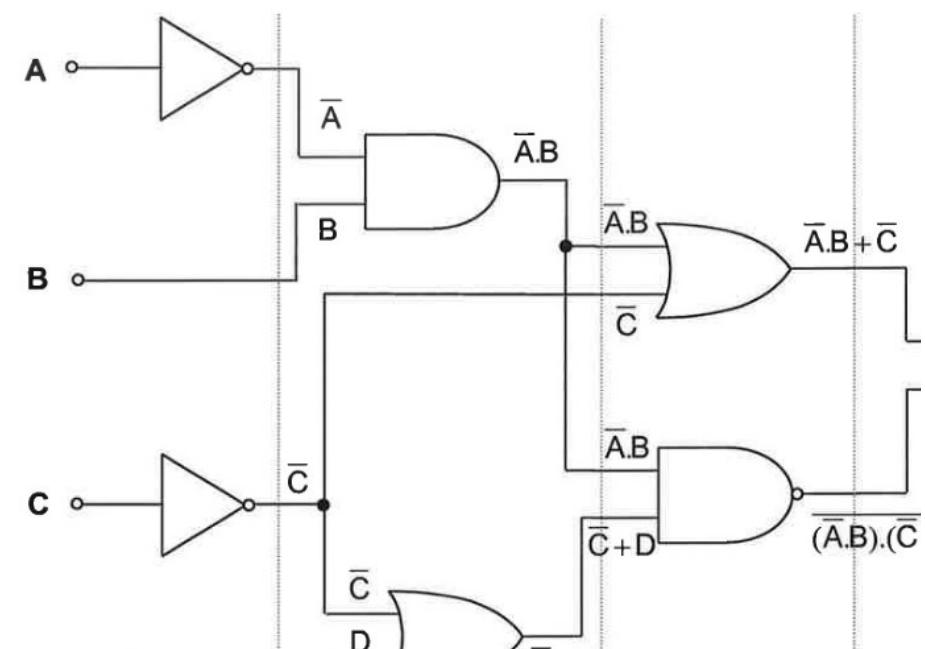
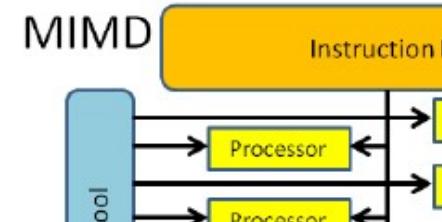
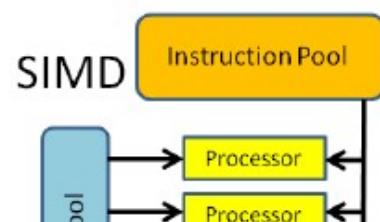
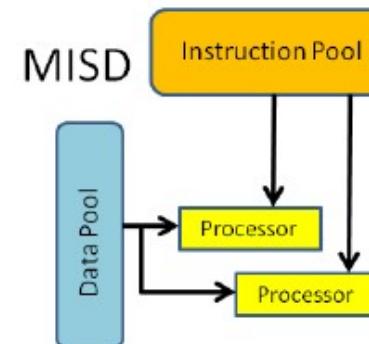
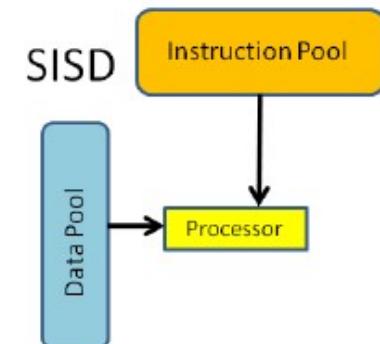
- Classification
- Parallel Prefix
- Graph Reachability
- Carry-Lookahead
- Sorting Networks

Classification

- *fine-grained parallelism*
vs. *distributed computing*
- SISD/MISD/SIMD/MIMD
- PRAMs: CRCW/CREW/EREW



Here: algorithms, not programs ("abstraction")



Primitives and Cost Measures

Size = #binary *gates*

Alternative **Size**

= #cores/#CPUs/#PCs,

→ Communication **Volume**,

Work := **Span** · #CPUs

Brent's Principle:
Work, Size ≥ Time

(seq.opt.) **Time**

Speedup := Time/Depth

Overhead := Size/Time

Depth = "parallel runtime"
Span

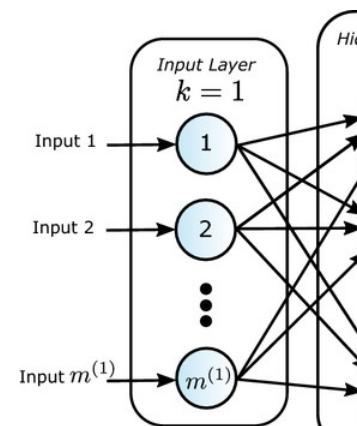
Here: Circuits

Lemma:

Any circuit of
binary gates,
depending on

N inputs, has

- size $\geq N-1$ and
- depth $\geq \log(N)$



Parallel N -ary Maximum

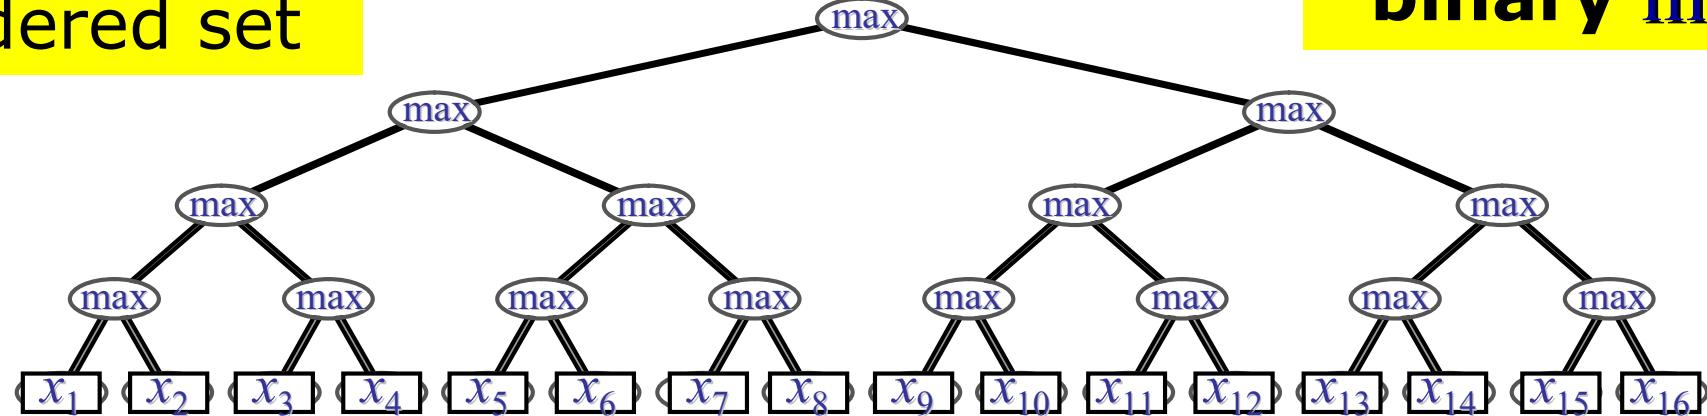
Size = $O(N)$ binary op.s

(X, \leq) totally ordered set

Depth = $O(\log N)$

optimal!

Primitive op:
binary max



Brent: **Size** \geq **Time**

(ignore *magnitude* of operands...)

(seq.opt.) **Time** = $O(N)$

Speedup = Time/Depth = $O(N/\log N)$

Overhead = Size/Time = $O(1)$

$(x_1, \dots, x_N) \rightarrow \max_{1 \leq n \leq N} x_n$

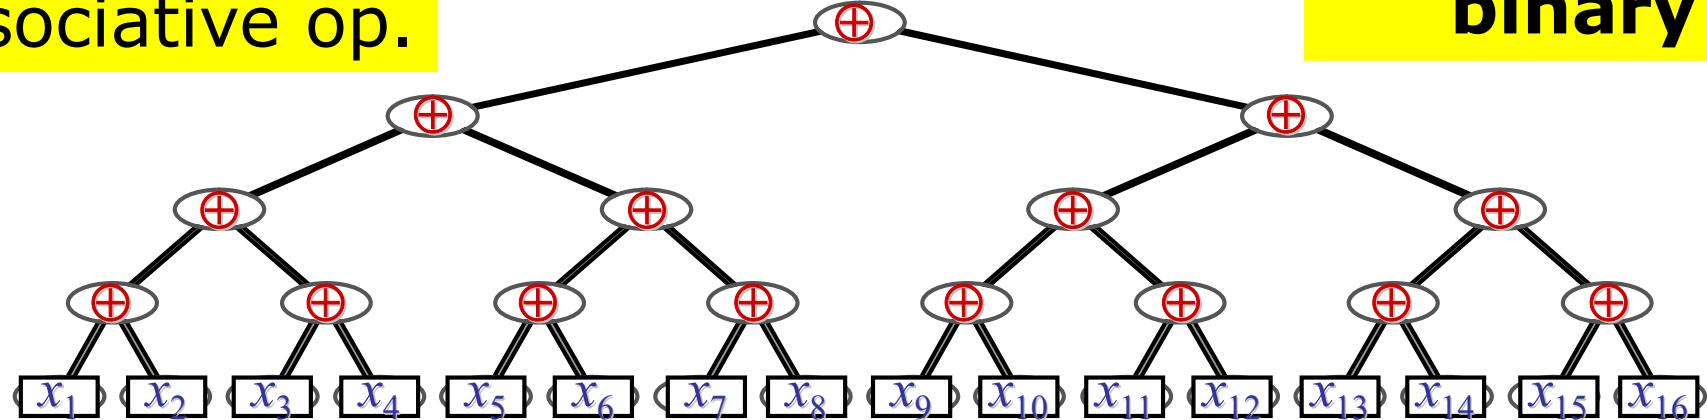
Parallel N -ary **associative** \oplus

Size = $O(N)$ binary op.s

(X, \oplus) set with
associative op.

Depth = $O(\log N)$

Primitive op:
binary \oplus



(seq.opt.) **Time** = $O(N)$

Speedup = Time/Depth = $O(N/\log N)$

Overhead = Size/Time = $O(1)$

$(x_1, \dots, x_N) \rightarrow$
 $\oplus_{1 \leq n \leq N} x_n$

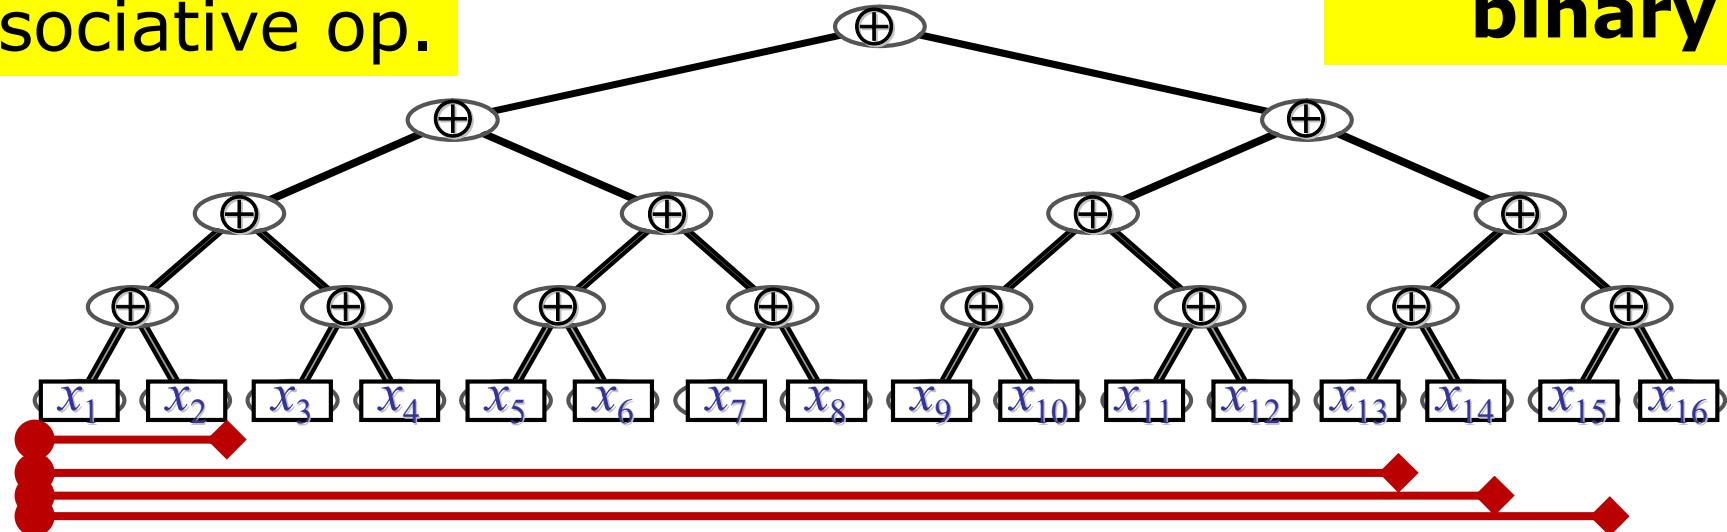
(Generalized) Prefix Sum / Scan

Size = $O(N^2)$ binary op.s

(X, \oplus) set with
associative op.

Depth = $O(\log N)$

Primitive op:
binary \oplus



(seq.opt.) **Time** = $O(N)$

Speedup = Time/Depth = $O(N/\log N)$

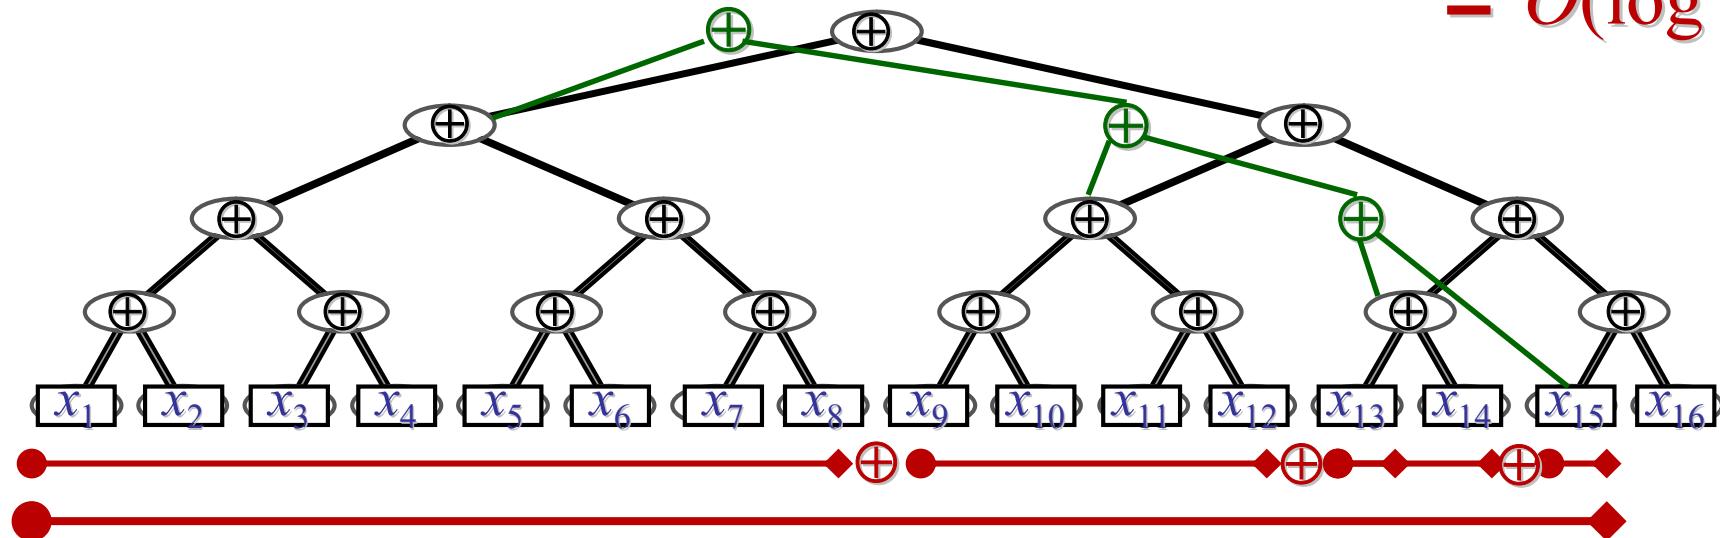
Overhead = Size/Time = $O(N)$

$(x_1, \dots, x_N) \rightarrow$
 $(\bigoplus_{1 \leq n \leq M} x_n : M=1 \dots N)$

Parallel Prefix

Size = $O(N)$ binary op.s
 $+ O(N) = O(N)$

Depth = $O(\log N)$
 $+ O(\log N)$
 $= O(\log N)$



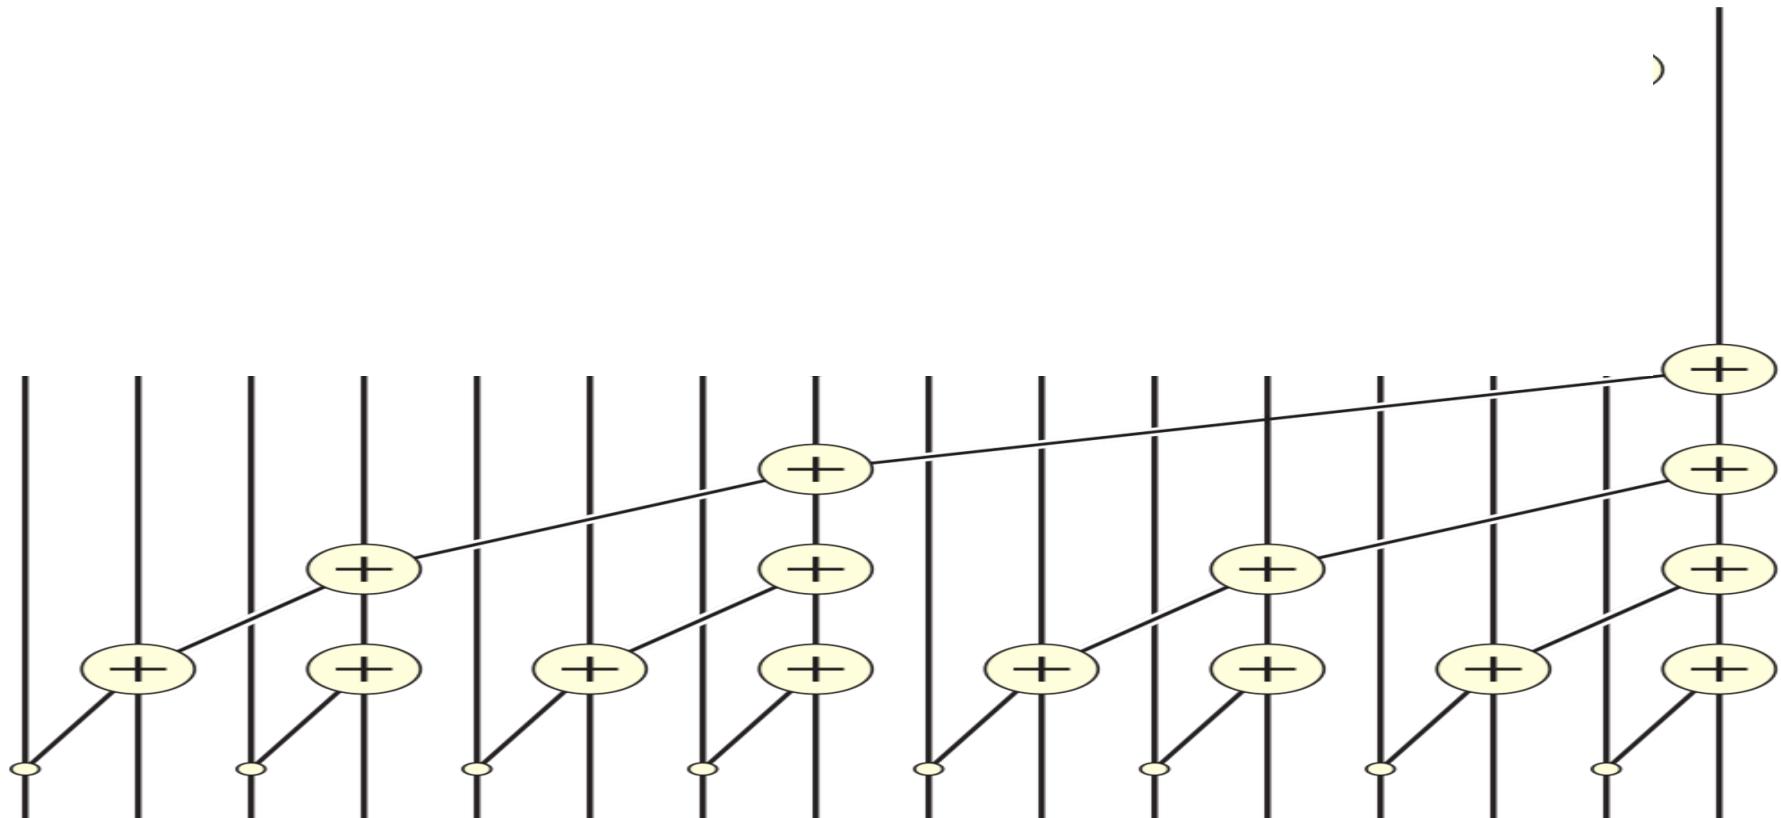
(seq.opt.) **Time** = $O(N)$

Speedup = Time/Depth = $O(N/\log N)$

Overhead = Size/Time = $O(1)$

$(x_1, \dots, x_N) \rightarrow$
 $(\bigoplus_{1 \leq n \leq M} x_n : M=1 \dots N)$

Parallel Prefix



(seq.opt.) **Time** = $O(N)$

Speedup = Time/Depth = $O(N/\log N)$

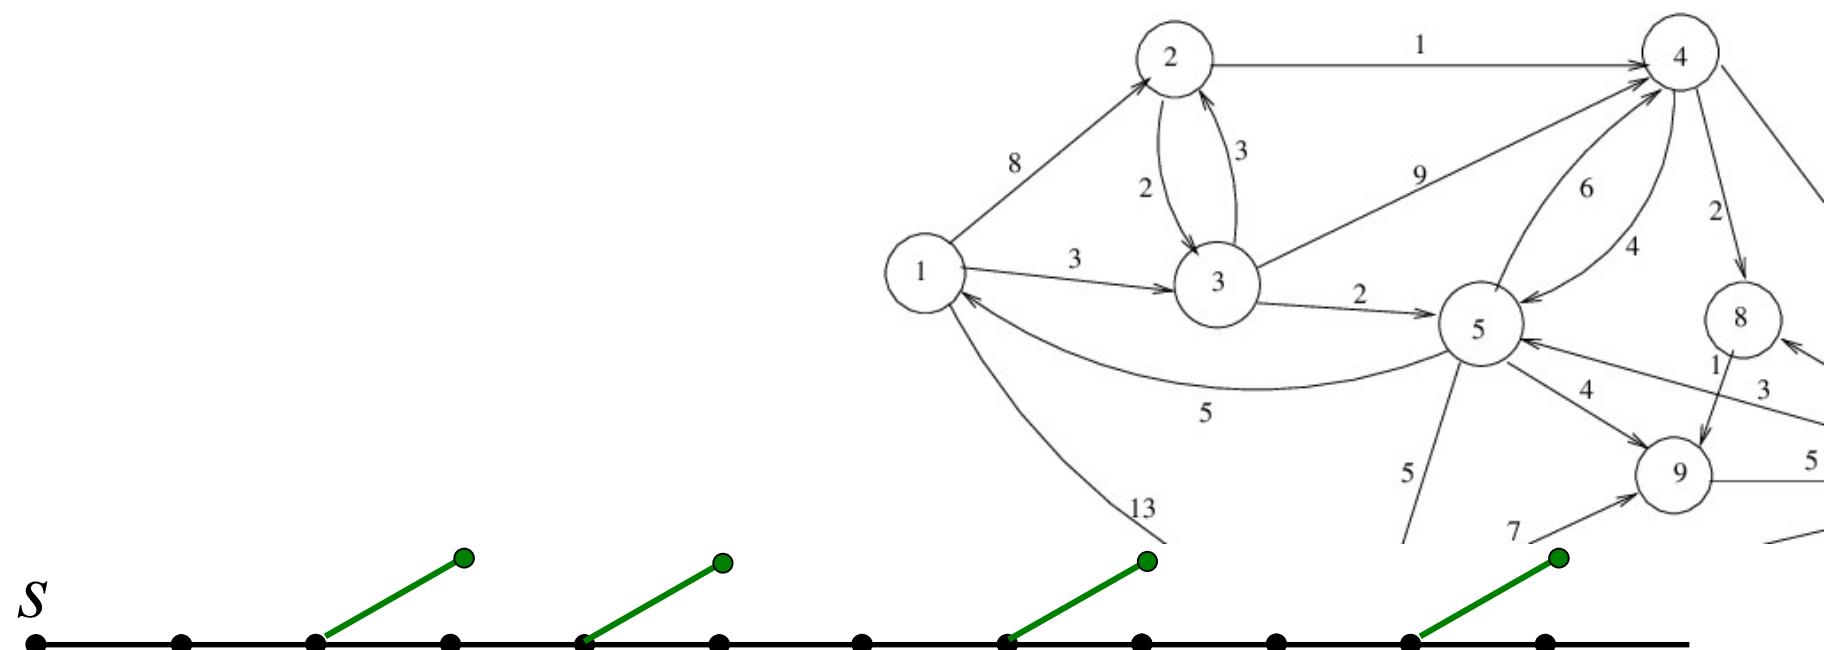
Overhead = Size/Time = $O(1)$

$$(x_1, \dots, x_N) \rightarrow (\bigoplus_{1 \leq n \leq M} x_n : M=1 \dots N)$$

Sequential Graph Reachability

Input: $K, s, t \in \mathbb{N}$ and directed graph as $N \times N$
0/1 adjacency matrix A , vertices $V = \{1 \dots N\}$

Output: Is t reachable from s within distance K ?



Dijkstra: Time = $O(N^2)$

Parallel Graph Reachability

Design & Analysis
of Algorithms
Martin Ziegler

Size = $O(N^3 \cdot \log K)$

Depth = $O(\log N \cdot \log K)$

Input: $K, s, t \in \mathbb{N}$ and directed graph as $N \times N$
0/1 adjacency matrix A , vertices $V = \{1 \dots N\}$

Output: Is t reachable from s within distance K ?

Answer: $(A^K)_{s,t}$ Bool. matrix power depth $O(\log N)$

$N \times N$ Boolean matrix product size $O(N^3)$

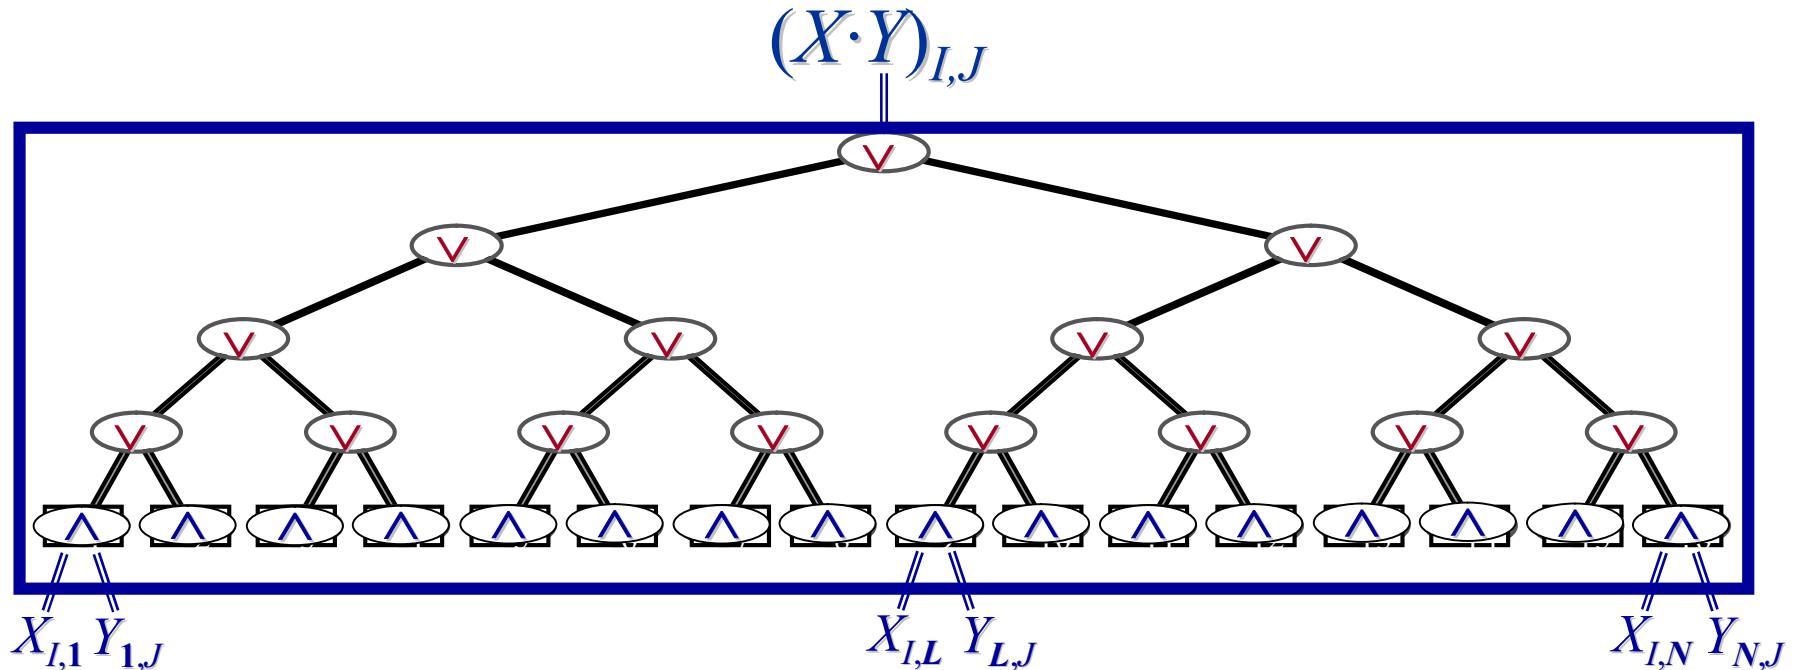
$$(X \cdot Y)_{I,J} = \bigvee_L (X_{I,L} \wedge Y_{L,J}) \quad \forall I, J = 1..N$$

Matrix powering $A \rightarrow A^K$
by repeated squaring A^2, A^4, A^8, \dots $O(\log K)$ repetitions

Dijkstra: Time = $O(N^2)$

Associative 

Explicit Parallel Reachability



$$(X \cdot Y)_{I,J} = \bigvee_L X_{I,L} \wedge Y_{L,J} \quad \forall I, J = 1..N$$

Matrix powering $A \rightarrow A^K$
by repeated squaring A^2, A^4, A^8, \dots $O(\log K)$ repetitions

Explicit Parallel Reachability

Design & Analysis
of Algorithms
Martin Ziegler

$$(A \cdot A)_{1,1}$$

.....

$$(A \cdot A)_{I,J}$$

.....

$$(A \cdot A)_{N,N}$$

$$A \cdot A = A^2$$

$$(A^2 \cdot A^2)_{1,1}$$

.....

$$(A^2 \cdot A^2)_{I,J}$$

.....

$$(A^2 \cdot A^2)_{N,N}$$

$$A^2 \cdot A^2 = A^4$$

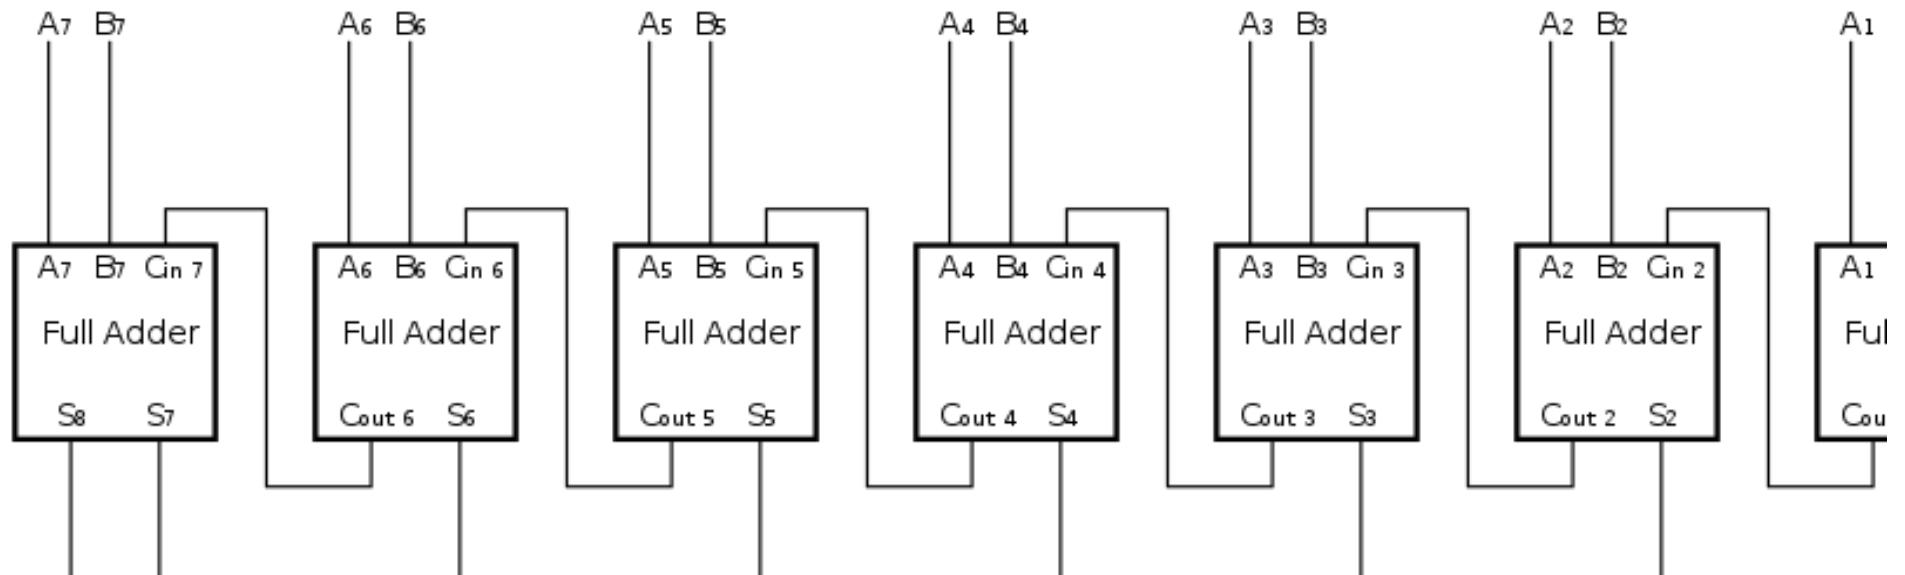
$$\vdots \quad \quad \quad (X \cdot Y)_{I,J} = \bigvee_L X_{I,L} \wedge Y_{L,J} \quad \quad \quad \forall I, J = 1..N \quad \vdots$$

Matrix powering $A \rightarrow A^K$
by repeated squaring A^2, A^4, A^8, \dots $O(\log K)$ repetitions

Sequential Ripple-Carry Adder

Input: two n -bit integers in binary,

$$A = (A_0, \dots, A_{N-1})_2 \quad \text{and} \quad B = (B_0, \dots, B_{N-1})_2$$



(seq.opt.) **Time** = $O(N)$

Speedup = Time/Depth = $O(1)$

Output: $S := A + B$ in binary,

$$S = (S_0, \dots, S_{N-1}, S_N)_2$$

Depth = $O(N)$

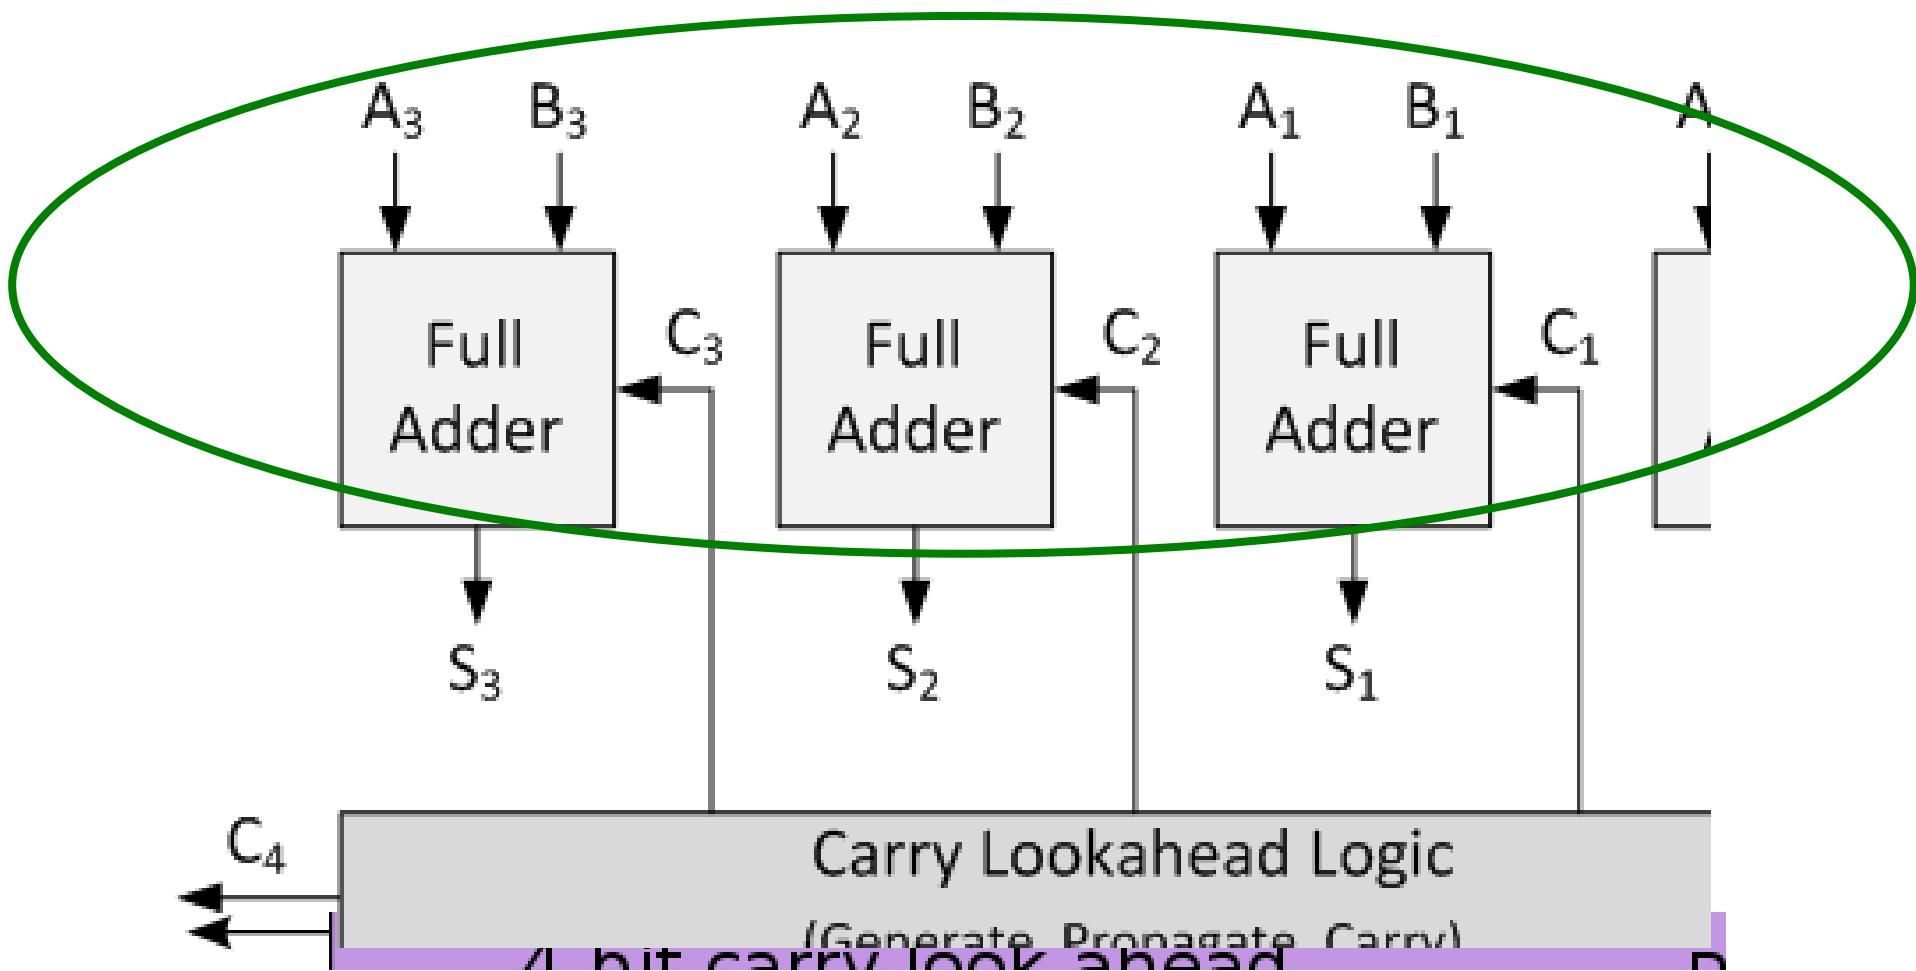
Operations:

binary \wedge, \vee, \neg

Carry-Lookahead Adder

Size = $O(N)$ binary bit op.s
 $+ O(N) = O(N)$

Depth = $O(\log N)$
 $+ O(1) = O(\log N)$



Carry Predictor

Size = $O(N^3)$

Depth = $O(\log N)$

$g_M := A_M \wedge B_M$ **generate** new carry to stage # $M+1$

$p_M := A_M \vee B_M$ **propagate** (possible) previous carry

$$C_{M+1} := g_M \vee (p_M \wedge C_M)$$

$$C_N = g_{N-1} \textcircled{V} g_{N-2} \wedge p_{N-1} \textcircled{V} \textcircled{g}_{N-3} \wedge p_{N-2} \wedge p_{N-1} \textcircled{V} \dots$$

$$\dots \textcircled{V} g_0 \wedge p_1 \wedge p_2 \wedge \dots \wedge p_{N-2} \wedge p_{N-1}$$

Final Goal: Calculate $(C_1 \dots C_N)$ in *parallel!*

Carry Predictor /Prefix Sum

Design & Analysis
of Algorithms
Martin Ziegler

Size = $O(N)$

Depth = $O(\log N)$

$g_M := A_M \wedge B_M$ generate new carry to stage # $M+1$

$p_M := A_M \vee B_M$ propagate (possible) previous carry

$$C_{M+1} := g_M \vee (p_M \wedge C_M) \quad P_{M+1} := p_M \wedge P_M \quad P_0 = 0$$

Lemma: The following operation \textcircled{C} on pairs of bits

$$(C, P) \textcircled{C} (C', P') := (C \vee (P \wedge C'), P \wedge P')$$

is (a) *idempotent* and (b) *associative*.

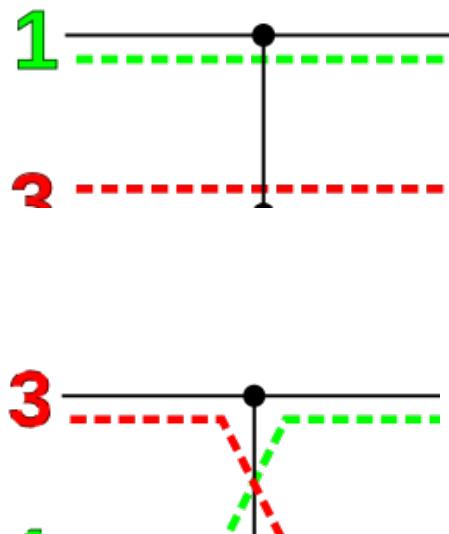
Power of Abstraction!

$(x_1, \dots, x_N) \rightarrow (\textcircled{C}_{1 \leq n \leq M} x_n : M=1\dots N), \textcircled{C}$ associative

Parallel Sorting, Revisited

(X, \leq) totally ordered set

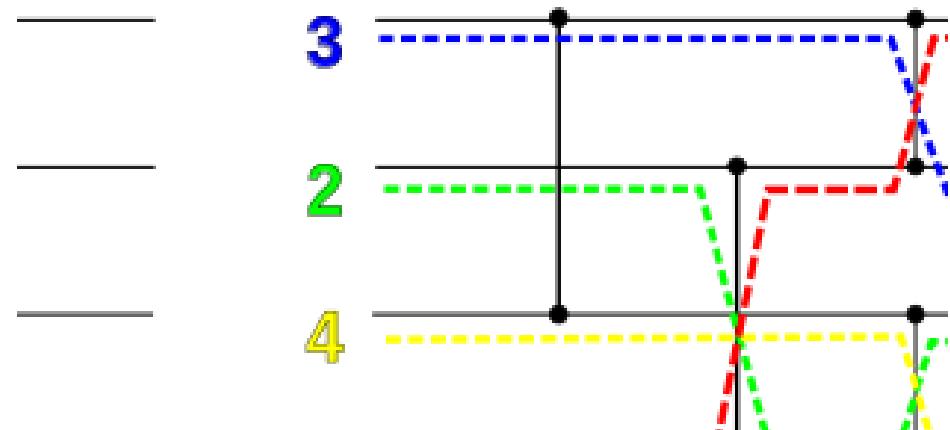
primitive/gate:



Gate semantics ($n > m$):

"**If** $x[n] < x[m]$ **then**
swap $x[n] \leftrightarrow x[m]$ "

Example: Sort4

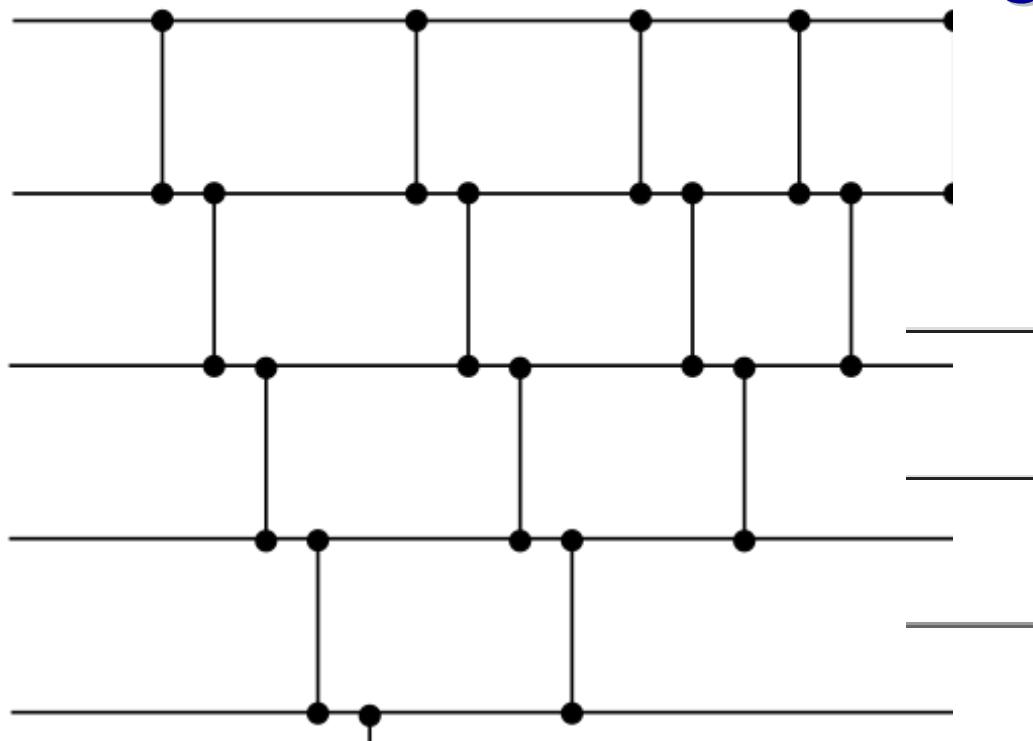


Bubble Sorting Network

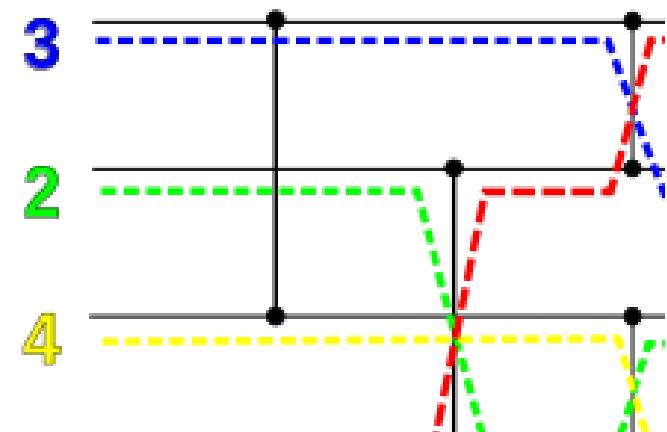
Size → opt.time $O(N \cdot \log N)$?

(Brent's Principle)

Depth → $O(\log N)$?



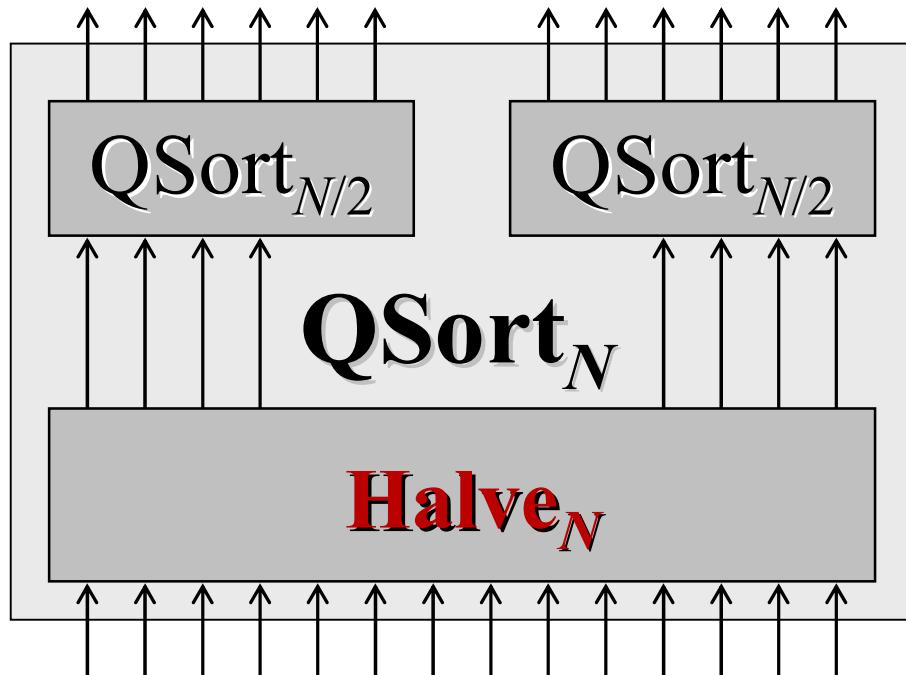
Bubble Sorting Network:
Size $O(N^2)$, Depth $O(N)$



Parallelizing *QuickSort* ?

Size → opt.time $O(N \cdot \log N)$?

Depth → $O(\log N)$?

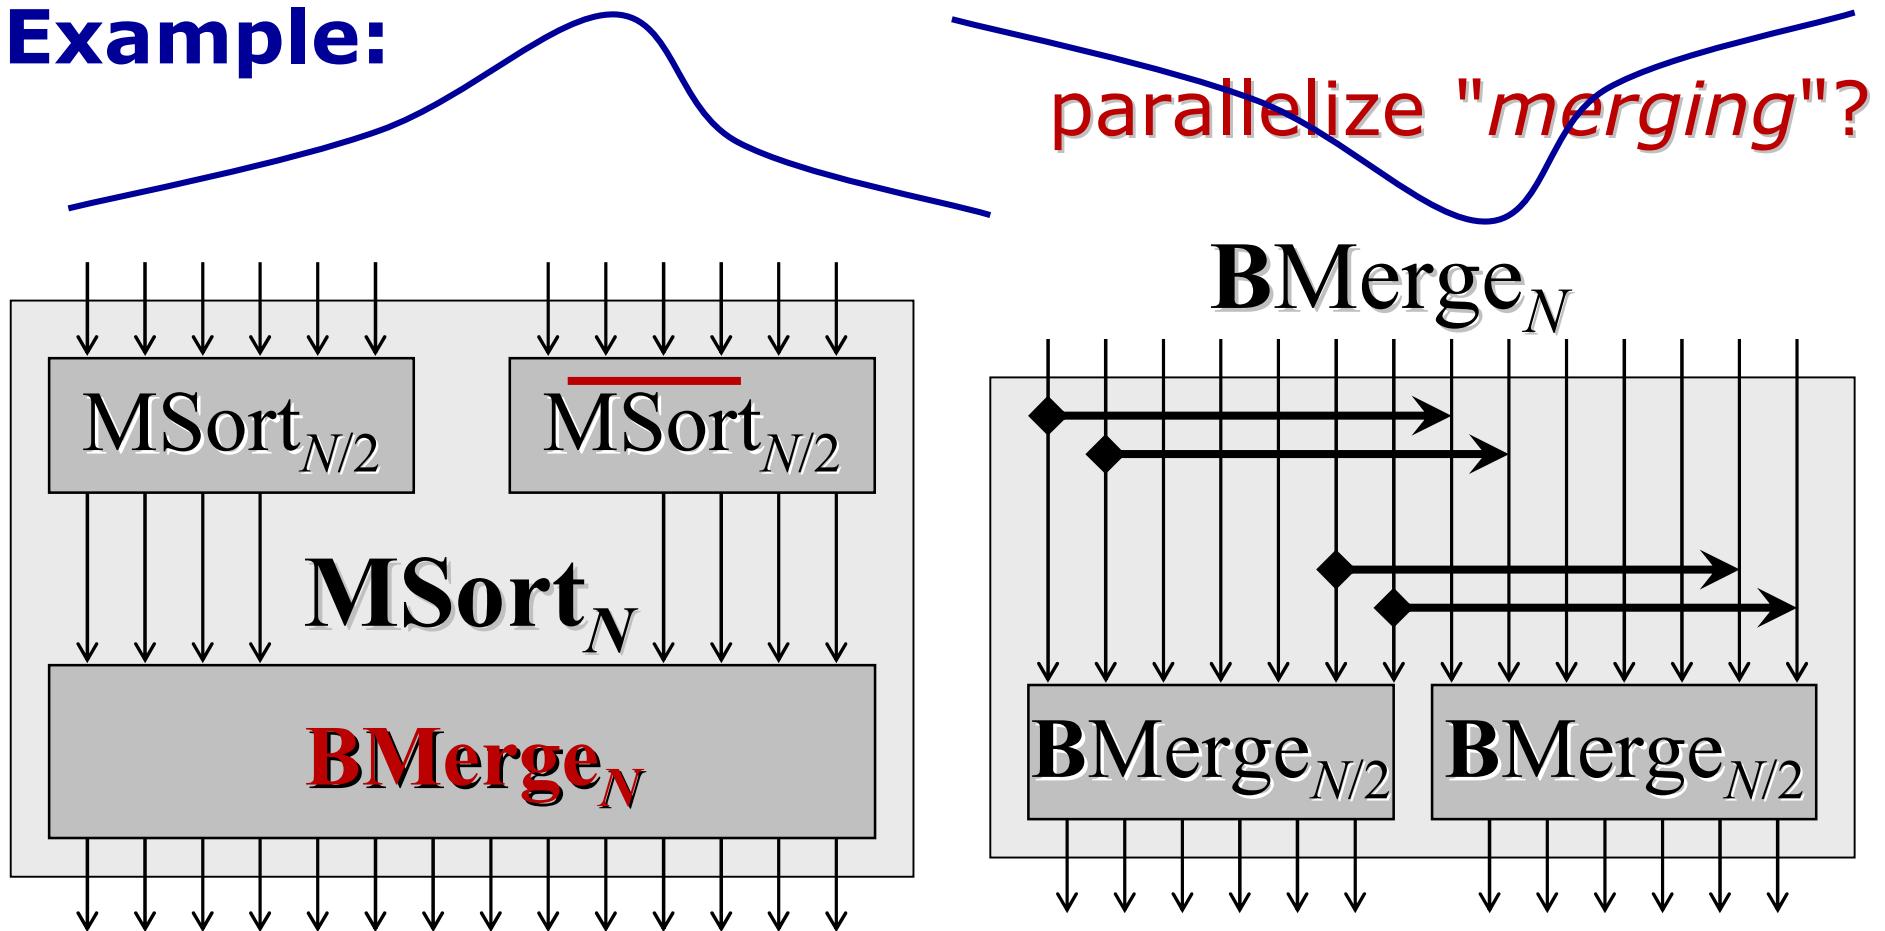


Bubble Sorting Network:
Size $O(N^2)$, Depth $O(N)$

Recall: *QuickSort* requires smart "halving" !

Parallelizing Mergesort ?

Example:

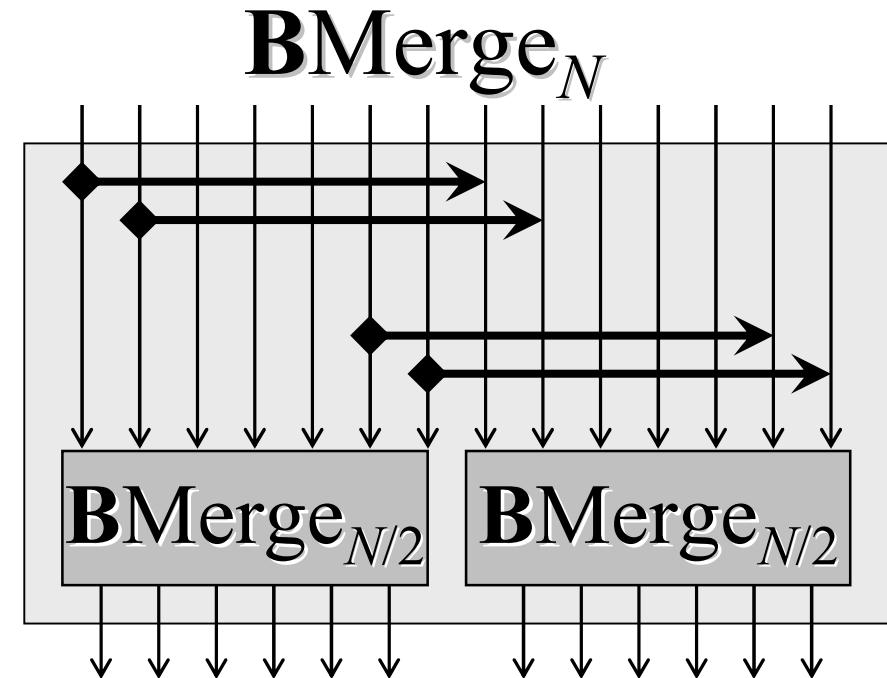
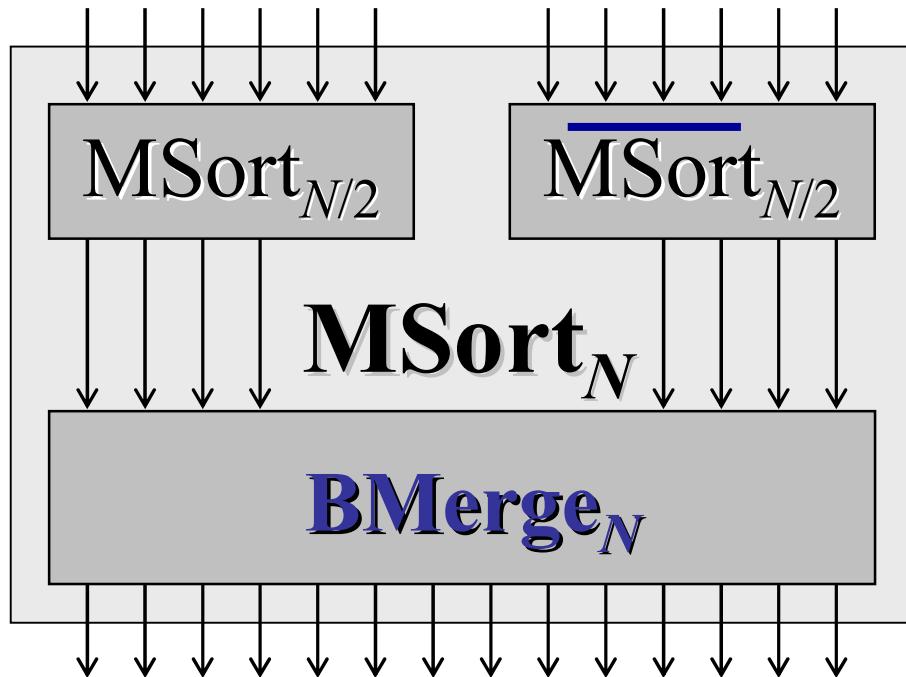


Def: Call $x[1\dots N]$ **bitonic** if, for some $M \leq N$,
 $x[1\dots M]$ non-decreasing & $x[M+1\dots N]$ non-increasing

or cyclic

Batcher's Bitonic Sorter

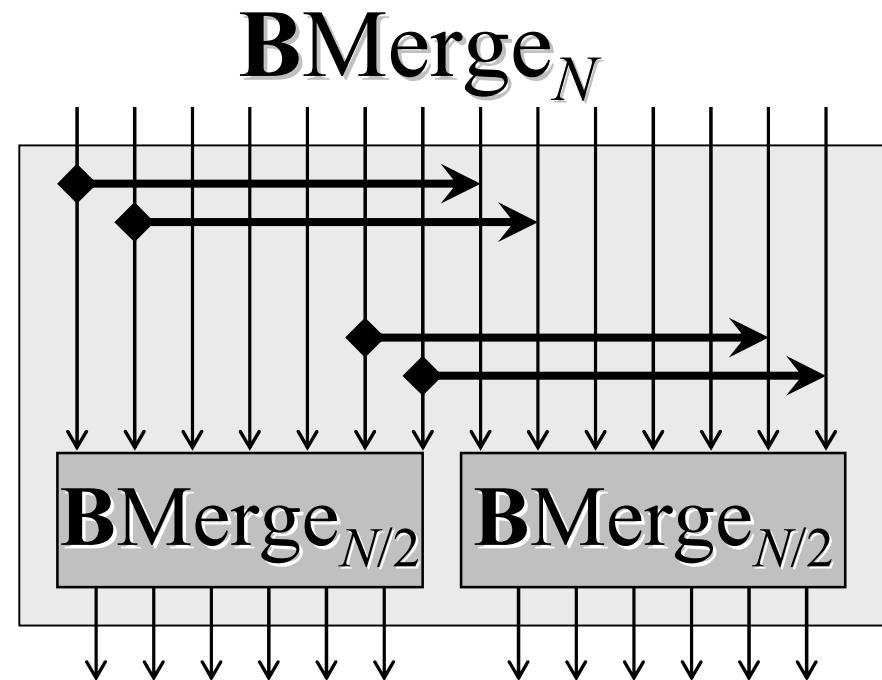
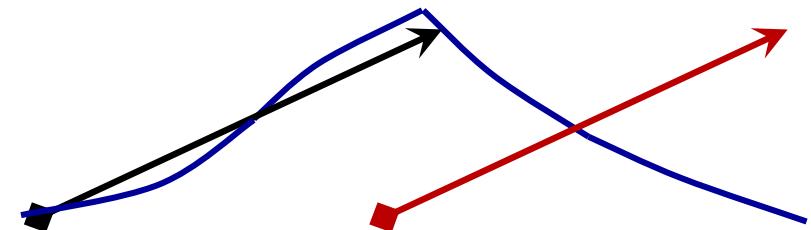
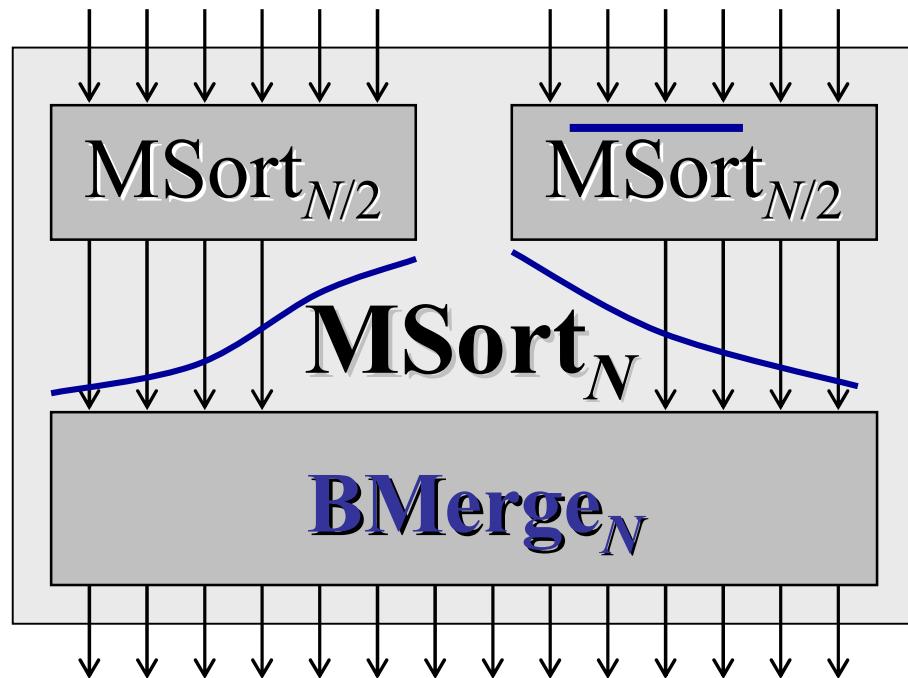
Size = $O(N \cdot \log^2 N)$ gates \rightarrow opt.time $O(N \cdot \log N)$?
Depth = $O(\log^2 N)$ par.time $\rightarrow O(\log N)$?



$$\text{DepthMSort}(N) = 1 \cdot \text{DepthMSort}(N/2) + \text{DepthBMerge}(N)$$

$$\text{DepthBMerge}(N) = O(1) + 1 \cdot \text{DepthBMerge}(N/2) = O(\log N)$$

Correctness of Batcher



or cyclic

Def: Call $x[1\dots N]$ **bitonic** if, for some $M \leq N$, ~~$M \leq N$~~ **shift**
 $x[1\dots M]$ non-decreasing & $x[M+1\dots N]$ non-increasing

AKS Sorting Network

Design & Analysis
of Algorithms
Martin Ziegler

Batcher's Bitonic Sorting Network

Size = $O(N \cdot \log^2 N)$ gates → opt.time $O(N \cdot \log N)$?

Depth = $O(\log^2 N)$ par.time → $O(\log N)$?

Ajtai, Komlós, Szemerédi (1983):

Sorting network of

Size $O(N \cdot \log N)$

Depth $O(\log N)$

§9 Summary

- Classification
- Parallel Prefix
- Graph Reachability
- Carry-Lookahead
- Sorting Networks