Chapter 10
Parallel Algorithms

Algorithm Theory
WS 2015/16

Fabian Kuhn

Sequential Algorithms

Classical Algorithm Design:

- One machine/CPU/process/… doing a computation

RAM (Random Access Machine):

- Basic standard model
- Unit cost basic operations
- Unit cost access to all memory cells

Sequential Algorithm / Program:

- Sequence of operations
  (executed one after the other)

Parallel and Distributed Algorithms

Today's computers/systems are not sequential:

- Even cell phones have several cores
- Future systems will be highly parallel on many levels
- This also requires appropriate algorithmic techniques

Goals, Scenarios, Challenges:

- Exploit parallelism to speed up computations
- Shared resources such as memory, bandwidth, …
- Increase reliability by adding redundancy
- Solve tasks in inherently decentralized environments
- …

Parallel and Distributed Systems

- Many different forms

- Processors/computers/machines/… communicate and share data through
  - Shared memory or message passing
- Computation and communication can be
  - Synchronous or asynchronous
- Many possible topologies for message passing
- Depending on system, various types of faults

Challenges

Algorithmic and theoretical challenges:

- How to parallelize computations
- Scheduling (which machine does what)
- Load balancing
- Fault tolerance
- Coordination / consistency
- Decentralized state
- Asynchrony
- Bounded bandwidth / properties of comm. channels
- …

Models

- A large variety of models, e.g.:

- **PRAM** (Parallel Random Access Machine)
  - Classical model for parallel computations
- **Shared Memory**
  - Classical model to study coordination / agreement problems, distributed data structures, …
- **Message Passing** (fully connected topology)
  - Closely related to shared memory models
- Message Passing in **Networks**
  - Decentralized computations, large parallel machines, comes in various flavors…

PRAM

- Parallel version of RAM model
- p processors, shared random access memory

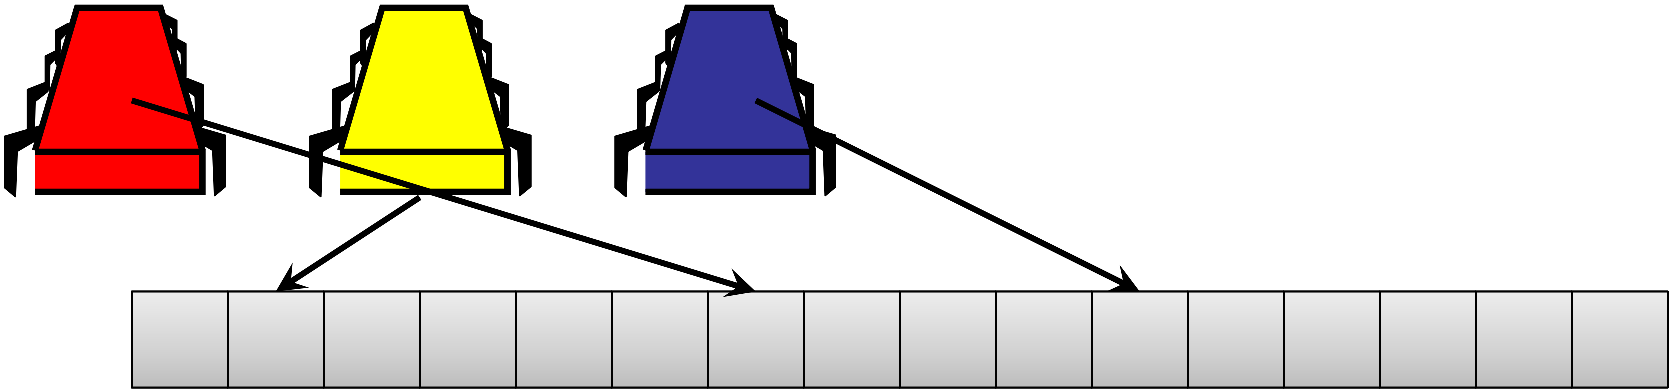

- Basic operations / access to shared memory cost 1
- Processor operations are synchronized time divided into rounds
- Focus on parallelizing computation rather than cost of communication, locality, faults, asynchrony, …

Other Parallel Models

- **Message passing:** Fully connected network, local memory and information exchange using messages

- **Dynamic Multithreaded Algorithms:** Simple parallel programming paradigm
  - E.g., used in Cormen, Leiserson, Rivest, Stein (CLRS)

```
FIB(n)
if n < 2
    then return n
x ← spawn FIB(n − 1)
y ← spawn FIB(n − 2)
sync
return (x + y)
```

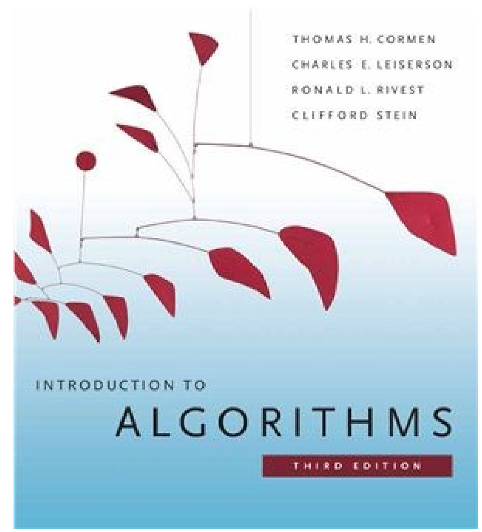

Parallel Computations

Sequential Computation:

- Sequence of operations

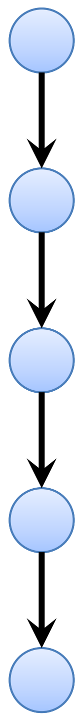

Parallel Computation:

- Directed Acyclic Graph (DAG)

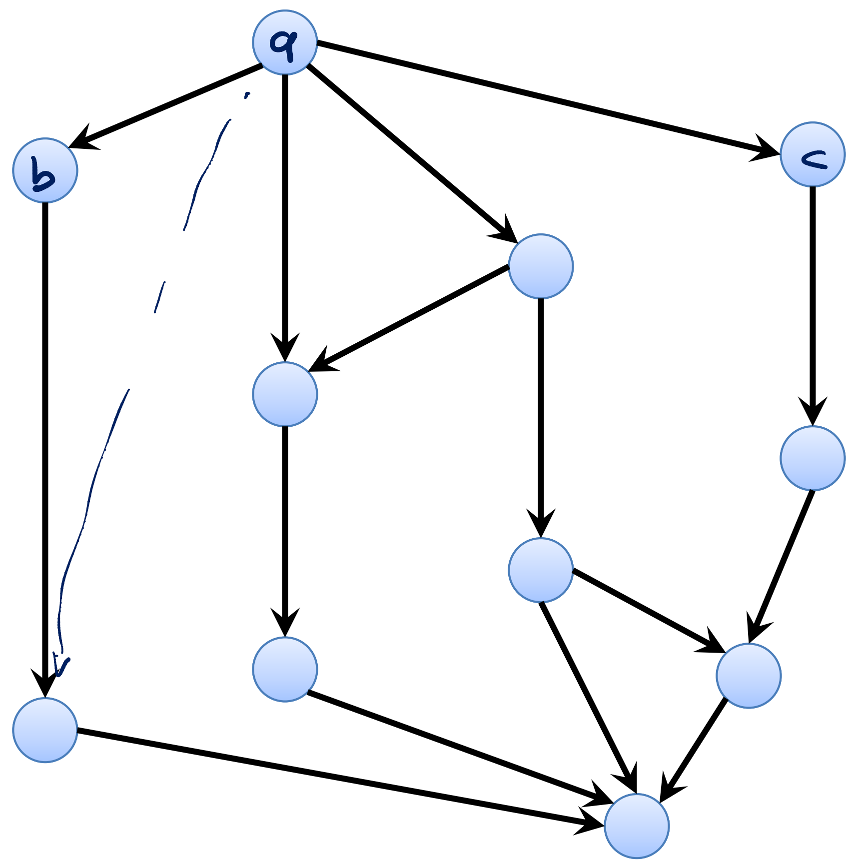

Parallel Computations

T_p: time to perform comp. with p procs

- T_1: **work** (total # operations)
  - Time when doing the computation sequentially
- T_∞: **critical path / span**
  - Time when parallelizing as much as possible
- **Lower Bounds**:

$$T_p \geq \left\lceil \frac{T_1}{p} \right\rceil, \qquad T_p \geq T_\infty$$

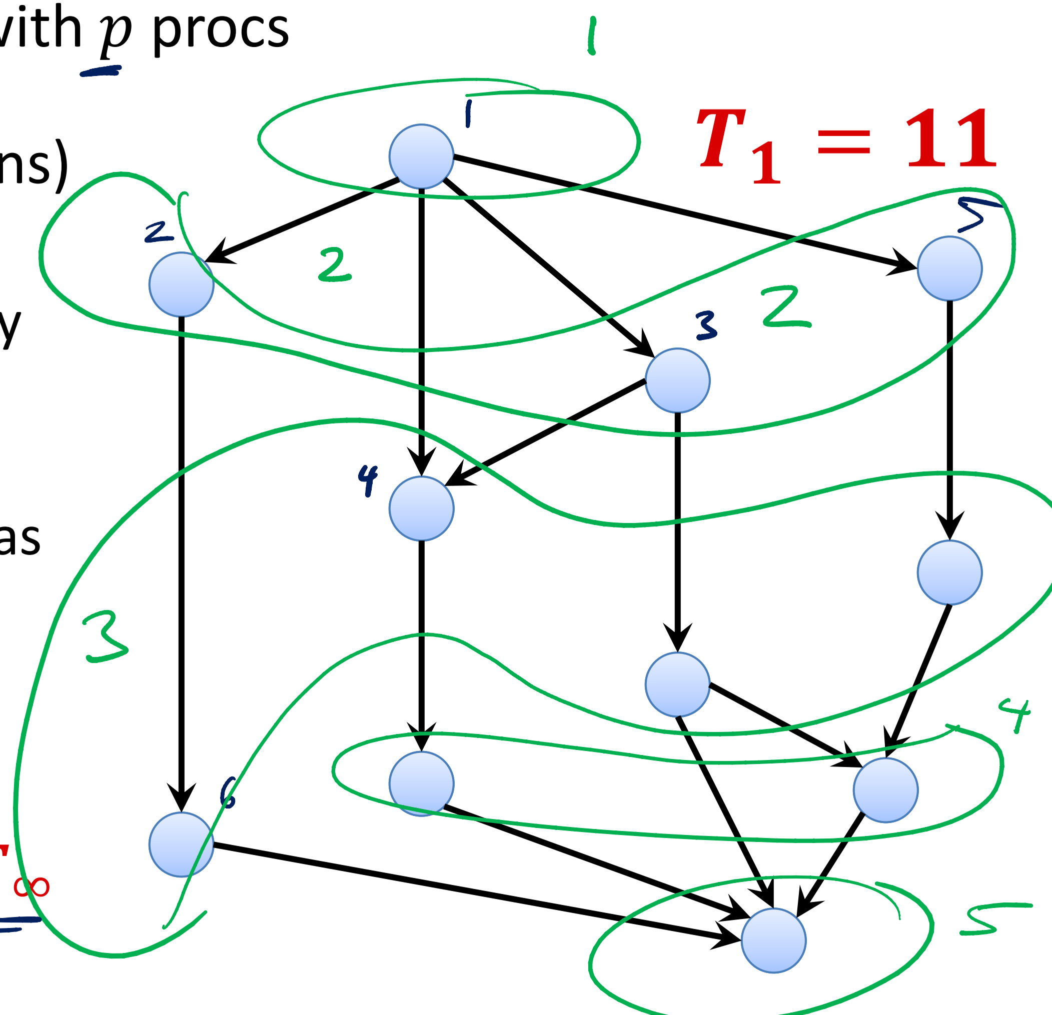

$T_1 = 11$

$T_\infty = 5$

Parallel Computations

$\boldsymbol{T_p}$: time to perform comp. with p procs

- **Lower Bounds**:

$$T_p \geq \frac{T_1}{p}, \qquad T_p \geq T_\infty$$

- **Parallelism**: $\dfrac{T_1}{T_\infty}$ ← one proc
  - maximum possible speed-up
- **Linear Speed-up**:

$$\frac{T_1}{T_p} = \Theta(p)$$

$\boldsymbol{T_1 = 11}$

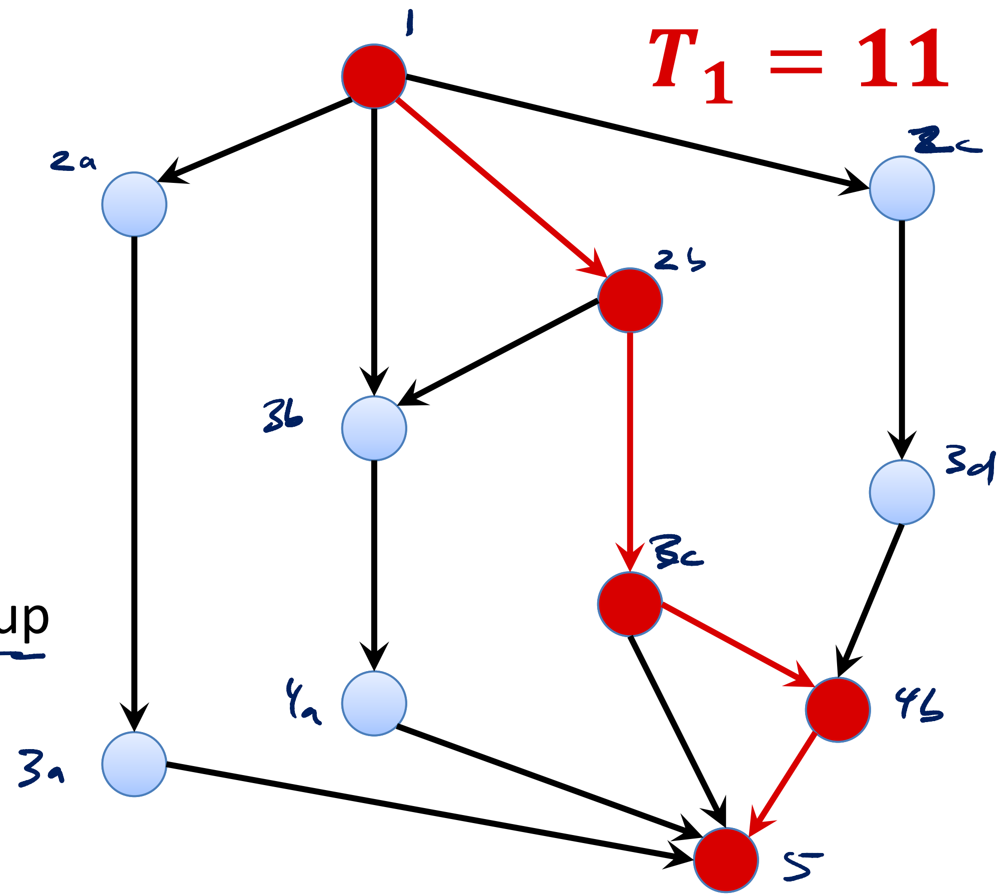

$\boldsymbol{T_\infty = 5}$

Scheduling

p procs.

T_1 ops.

- How to assign operations to processors?
- Generally an online problem
  - When scheduling some jobs/operations, we do not know how the computation evolves over time

Greedy (offline) scheduling:

- Order jobs/operations as they would be scheduled optimally with ∞ processors (topological sort of DAG)
  - Easy to determine: With ∞ processors, one always schedules all jobs/ops that can be scheduled
- Always schedule as many jobs/ops as possible
- Schedule jobs/ops in the same order as with ∞ processors
  - i.e., jobs that become available earlier have priority

Brent's Theorem

Brent's Theorem: On p processors, a parallel computation can be performed in time

$$T_p \le \frac{T_1 - T_\infty}{p} + T_\infty. \le \frac{T_1}{p} + T_\infty$$

Proof:

- Greedy scheduling achieves this…
- #operations scheduled with ∞ processors in round i: x_i

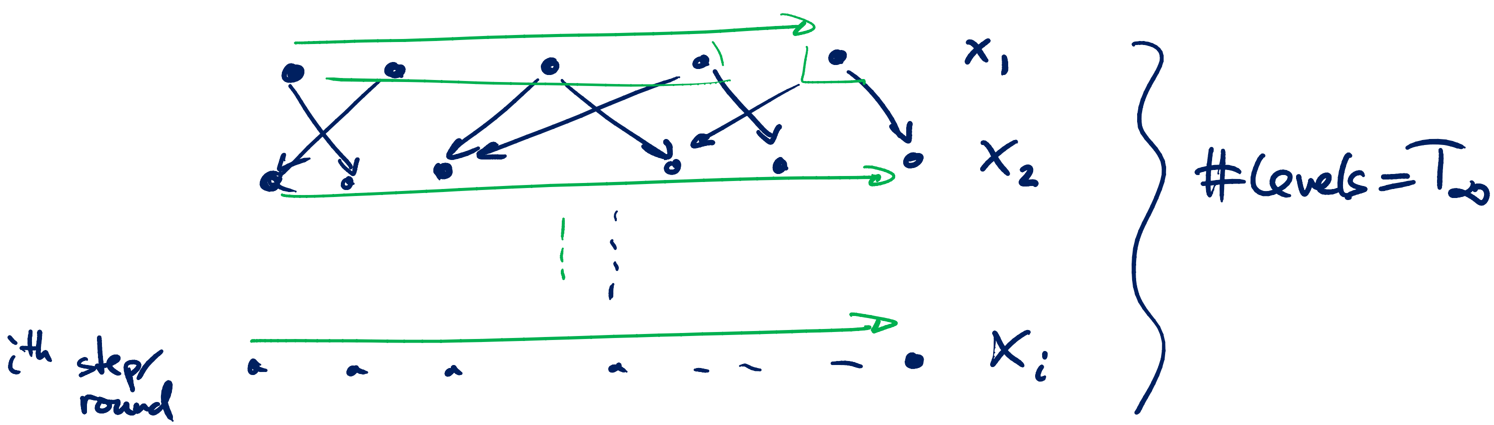

Brent's Theorem

Brent's Theorem: On p processors, a parallel computation can be performed in time

$$T_p \leq \frac{T_1 - T_\infty}{p} + T_\infty.$$

Proof:

- Greedy scheduling achieves this…
- #operations scheduled with ∞ processors in round i: x_i

p procs.

t_i : time to schedule x_i ops.

$$\sum_{i=1}^{T_\infty} x_i = T_1$$

$$t_i = \left\lceil \frac{x_i}{p} \right\rceil \leq \frac{x_i}{p} + \frac{p-1}{p} = \frac{x_i - 1}{p} + 1$$

$$T_p^G \leq \sum_{i=1}^{T_\infty} t_i \leq \underbrace{\sum_{i=1}^{T_\infty} \frac{x_i}{p}}_{T_1/p} - \underbrace{\sum_{i=1}^{T_\infty} \frac{1}{p}}_{T_\infty/p} + \underbrace{\sum_{i=1}^{T_\infty} 1}_{T_\infty} = \frac{T_1 - T_\infty}{p} + T_\infty \leq \frac{T_1}{p} + T_\infty$$

Brent's Theorem

Brent's Theorem: On p processors, a parallel computation can be performed in time

$$T_p \le \underbrace{\frac{T_1 - T_\infty}{p}}_{\varphi} + T_\infty.$$

Corollary: Greedy is a 2-approximation algorithm for scheduling.

lower bounds

$T_p^* \geq \frac{T_1}{p}$

$T_p^* \geq T_\infty$

$$T_p^G \le \underbrace{\frac{T_1}{p}}_{\le T_p^*} + \underbrace{T_\infty}_{\le T_p^*} \le 2T_p^*$$

Corollary: As long as the number of processors $p = O(T_1/T_\infty)$, it is possible to achieve a linear speed-up.

PRAM

EREW PRAM

Back to the PRAM:

- Shared random access memory, synchronous computation steps
- The PRAM model comes in variants…

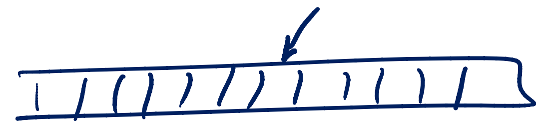

EREW (exclusive read, exclusive write):

to the same cell

- Concurrent memory access by multiple processors is not allowed
- If two or more processors try to read from or write to the same memory cell concurrently, the behavior is not specified

CREW (concurrent read, exclusive write):

- Reading the same memory cell concurrently is OK
- Two concurrent writes to the same cell lead to unspecified behavior

  also concurrent write & read not allowed

- This is the first variant that was considered (already in the 70s)

PRAM

The PRAM model comes in variants… EREW

CRCW (concurrent read, concurrent write):

- Concurrent reads and writes are both OK
- Behavior of concurrent writes has to specified
  - Weak CRCW: concurrent write only OK if all processors write 0
  - Common-mode CRCW: all processors need to write the same value
  - Arbitrary-winner CRCW: adversary picks one of the values
  - Priority CRCW: value of processor with highest ID is written
  - Strong CRCW: largest (or smallest) value is written

(ID, x)

- The given models are ordered in strength:

weak $\leq$ common-mode $\leq$ arbitrary-winner $\leq$ priority $\leq$ strong

Some Relations Between PRAM Models

Theorem: A parallel computation that can be performed in time t, using p proc. on a strong CRCW machine, can also be performed in time $O(t \log p)$ using p processors on an EREW machine.

- Each (parallel) step on the strong CRCW machine can be simulated by $O(\log p)$ steps on an EREW machine

concurrent writes: mem. cell i, coordinate write access to cell i

cell i : auxiliary mem. cells (initialized to 0)

say: $p = 8$

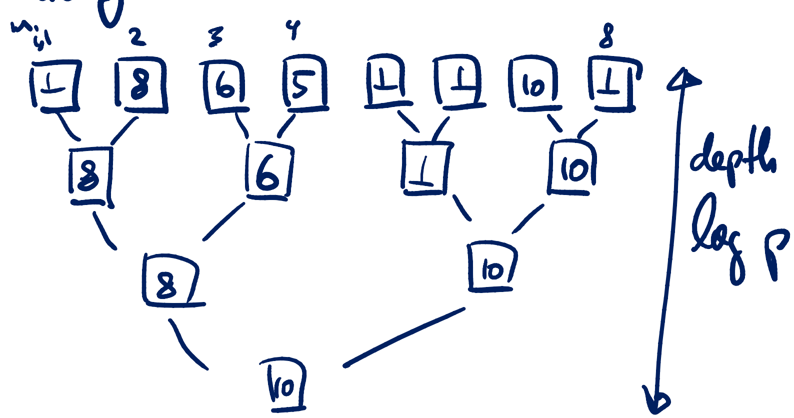

proc.	2	writes	8
	3	"	6
	4	"	5
	7	"	10

Some Relations Between PRAM Models

Theorem: A parallel computation that can be performed in time t, using p proc. on a strong CRCW machine, can also be performed in time $O(t \log p)$ using p processors on an EREW machine.

- Each (parallel) step on the CRCW machine can be simulated by $O(\log p)$ steps on an EREW machine

Concurrent reads

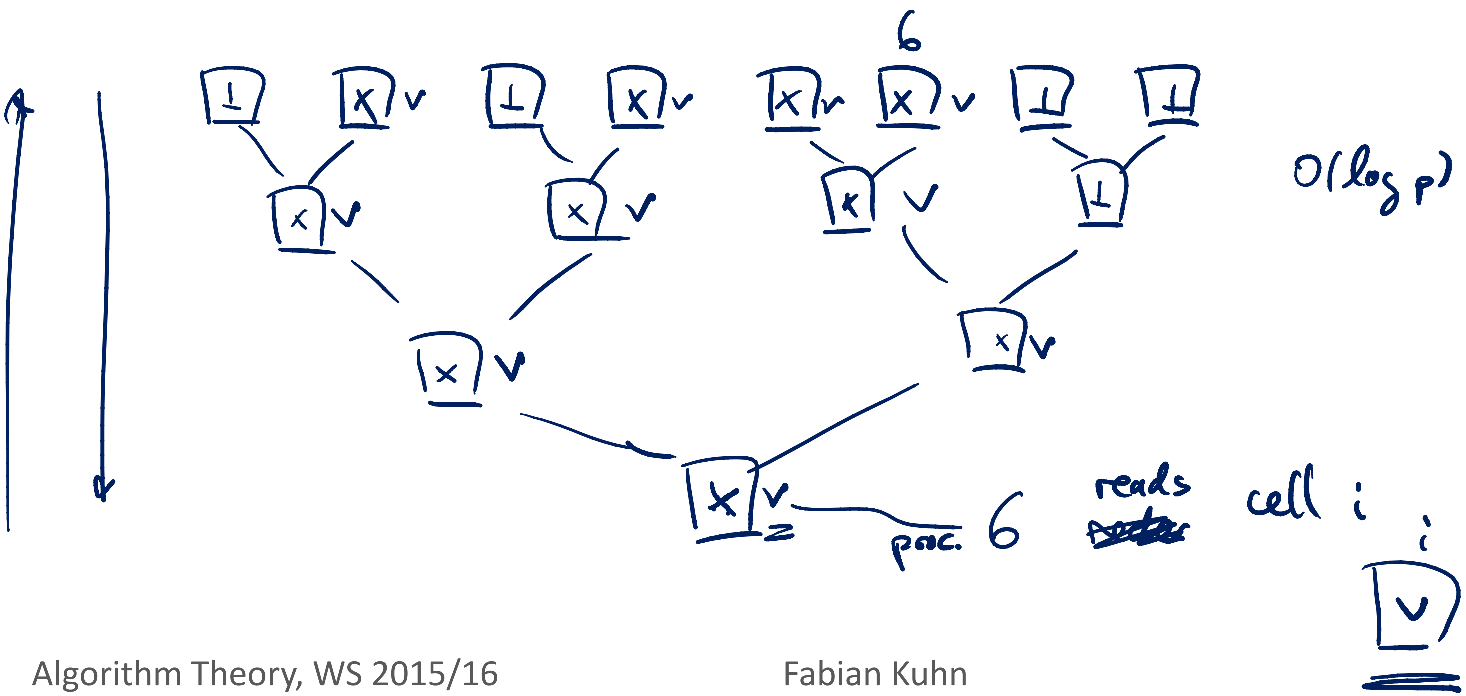

$O(\log p)$ parallel steps

Some Relations Between PRAM Models

Theorem: A parallel computation that can be performed in time t, using p proc. on a strong CRCW machine, can also be performed in time $O(t \log p)$ using p processors on an EREW machine.

- Each (parallel) step on the CRCW machine can be simulated by $O(\log p)$ steps on an EREW machine

Theorem: A parallel computation that can be performed in time t, using p probabilistic processors on a strong CRCW machine, can also be performed in expected time $O(t \log p)$ using $O(p/\log p)$ processors on an arbitrary-winner CRCW machine.

- The same simulation turns out more efficient in this case

Some Relations Between PRAM Models

Theorem: A computation that can be performed in time t, using p processors on a strong CRCW machine, can also be performed in time $O(t)$ using $O(p^2)$ processors on a weak CRCW machine

Proof:

- Strong: largest value wins, weak: only concurrently writing 0 is OK

Simulate 1 step of a strong CRCW PRAM on a weak CRCW PRAM

Processes: strong CRCW: $1, \dots, p$

additional procs: $q_{i,j}$ for every pair (i,j), $i,j \in \{1, \dots, p\}$

additional mem. cells:

for all $i \in \{1, \dots, p\}$: f_i, v_i, a_i (initialized to 0)

if proc. $i \in \{1, \dots, p\}$ wants to write x to mem. cell c (in strong CRCW)

$f_i := 1$, $v_i := x$, $a_i := c$

Some Relations Between PRAM Models

Theorem: A computation that can be performed in time t, using p processors on a strong CRCW machine, can also be performed in time $O(t)$ using $O(p^2)$ processors on a weak CRCW machine

Proof:

- Strong: largest value wins, weak: only concurrently writing 0 is OK

proc i wants to write x to cell c : $f_i := 1$, $v_i := x$, $a_i := c$

$\forall i,j$: $q_{i,j}$ reads $f_i, f_j, v_i, v_j, a_i, a_j$ (assume $i<j$)

if $f_i = f_j = 1$ and $a_i = a_j$ then

if $v_j \geq v_i$ then $f_i := 0$
else $f_j := 0$
} req. concurrent writes

procs

1	2	3	4	5
10	12	11	12	11

proc. i writes v_i to cell a_i $\Longleftrightarrow$ $f_i = 1$

f_i: $q_{1,2}$: $f_1 := 0$
$q_{2,3}$: $f_3 := 0$
$q_{2,4}$: $f_2 := 0$

Computing the Maximum

Given: n values

Goal: find the maximum value

Observation: The maximum can be computed in parallel by using a binary tree.

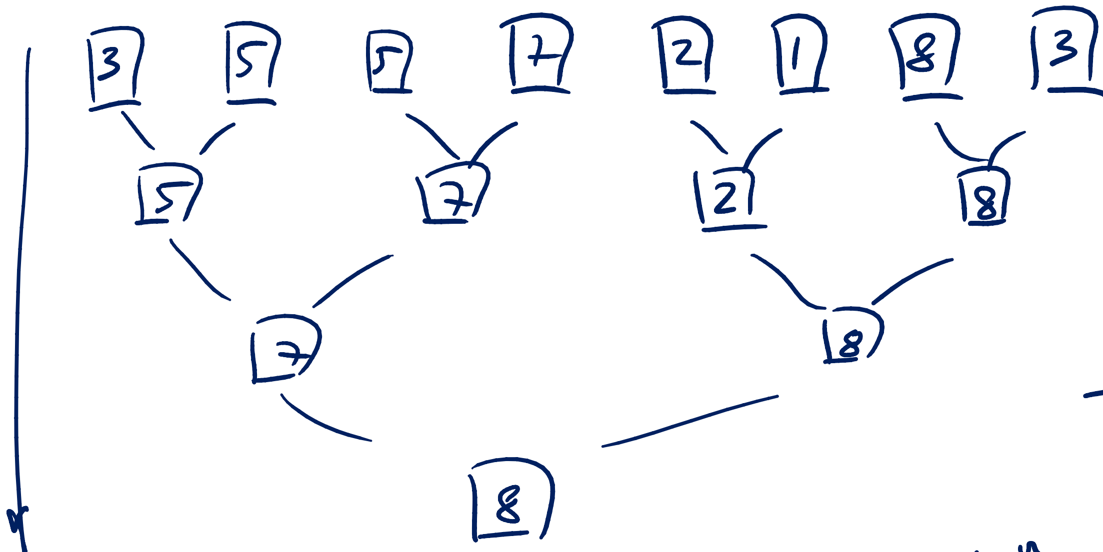

can be done on an EREW PRAM

$T_1 = O(n)$

$T_\infty = O(\log n)$

$T_p = O\left(\frac{n}{p} + \log n\right)$

Computing the Maximum

Observation: On a strong CRCW machine, the maximum of a n values can be computed in $O(1)$ time using n processors

- Each value is concurrently written to the same memory cell

Lemma: On a weak CRCW machine, the maximum of n integers between 1 and $\sqrt{n}$ can be computed in time $O(1)$ using $O(n)$ proc.

Proof:

- We have $\sqrt{n}$ memory cells $f_1, \dots, f_{\sqrt{n}}$ for the possible values
- Initialize all $f_i \coloneqq 1$
- For the n values $x_1, \dots, x_n$, processor j sets $f_{x_j} \coloneqq 0$
  - Since only zeroes are written, concurrent writes are OK
- Now, $f_i = 0$ iff value i occurs at least once
- Strong CRCW machine: max. value in time $O(1)$ w. $O(\sqrt{n})$ proc.
- Weak CRCW machine: time $O(1)$ using $O(n)$ proc. (prev. lemma)