## 캐zürich <br> Izurich

## Torsten Hoefler

## Parallel Programming Message Passing \& Parallel

 Sorting (A taste of parallel algorithms!)

EHzürich


\section*{User Lab Usage by Research Field <br> | Research Field | Node | \% |
| :---: | :---: | :---: |
| Physics | 15537867 | 35 |
| Chemistry \& Materials | 14370198 | 33 |
| Earth \& Environmental Science | 4736388 | 11 |
| Mechanics $\alpha$ Engineering | 3650101 | 8 |
| Life Science | 3308827 | 8 |
| Others | 2175327 | 5 |
| Total Usage | 43778708 | 100 |

## User Lab Usage by Institution




禺 Which programming languages and parallelization paradigms are you using primarily?


## Last week

## - Transactional memory

- Motivation (locks are bad, wait-/lock-free is hard)
- Concepts (Atomicity, Consistency, Isolation - ACI(D))
- Implementation options (keep track of read and write sets)
- Example: dining philosophers


## - Distributed memory

- Isolation of state - big simplification
- Event-driven messaging/Actors (example: Erlang)
- CSP-style (example: Go)


## Learning goals for today

- Finish Go


## - Message Passing Interface

- Standard library for high-performance parallel programming
- Processes, communicators, collectives - concepts of distributed memory programming
- Matching, deadlocks - potential pitfalls
- A primer on parallel algorithms
- Parallel sorting
- Sorting with fixed structures - sorting networks


## Example: Concurrent prime sieve

Each station removes multiples of the first element received and passes on the remaining elements to the next station


## Concurrent prime sieve

```
    func Generate(ch chan<- int) {
    for i := 2; ; i++ {
        ch <- i
    }
}
```

```
func Filter(in <-chan int, out chan<- int, prime int) {
    for {
        i :=<-in // Receive value from 'in'.
        if i%prime != 0 {
        out <- i // Send 'i' to 'out'.
        }
```

func main() {

```
    ch := make(chan int)
    go Generate(ch)
    for \(\mathbf{i}:=0\); \(\mathbf{i}<10\); i++ \{
        prime := <-ch

        fmt. Println(prime)
        ch1 := make(chan int)
        go Filter(ch, ch1, prime)
        ch \(=\) ch1
    \}
\}

\section*{Message Passing Interface (MPI)}


\section*{Message Passing Interface (MPI)}

Message passing libraries:
- PVM (Parallel Virtual Machines) 1980s
- MPI (Message Passing Interface) 1990s

\section*{MPI = Standard API}
- Hides Software/Hardware details
- Portable, flexible
- Implemented as a library


\section*{Process Identification}
- MPI processes can be collected into groups
- Each group can have multiple colors (some times called context)
- Group + color == communicator (it is like a name for the group)
- When an MPI application starts, the group of all processes is initially given a predefined name called MPI_COMM_WORLD
- The same group can have many names, but simple programs do not have to worry about multiple names
- A process is identified by a unique number within each communicator, called rank
- For two different communicators, the same process can have two different ranks: so the meaning of a "rank" is only defined when you specify the communicator

\section*{MPI Communicators}
- Defines the communication domain of a communication operation: set of processes that are allowed to communicate with each other.

- Initially all processes are in the communicator MPI_COMM_WORLD.

- The rank of processes are associated with (and unique within) a communicator, numbered from 0 to n -1

\section*{Communicators}

Communicators do not need to contain all processes in the system

Every process in a communicator has an ID called as "rank"

When you start an MPI program, there is one predefined communicator MPI_COMM_WORLD

Can make copies of this communicator (same group of processes, same ranks, but different "aliases")


The same process might have different ranks in different communicators

Communicators do not

Communicators can be created "by hand" or using tools
Simple programs typically only use the predefined communicator MPI_COMM_WORLD
(which is sometimes considered bad practice because of modularity issues)

\section*{Process Ranks}

Processes are identified by nonnegative integers, called ranks
\(p\) processes are numbered \(0,1,2, . . p-1\)
```

public static void main(String args []) throws Exception {
MPI.Init(args);
// Get total number of processes (p)
int size = MPI.COMM_WORLD.Size();
// Get rank of current process (in [0..p-1])
int rank = MPI.COMM_WORLD.Rank();
MPI.Finalize();
}

```

\section*{SPMD}

Single Program

\section*{Multiple Data}
(Multiple Instances)


\section*{Communication}
```

void Comm.Send(
Object buf,
int offset,
int count,
Datatype datatype, data type of items, must be explicitely specified
int dest, destination process id
int tag data id tag
communicator
pointer to data to be sent
from MPJ Spec
number of items to be sent
)

```


\section*{Parallel Sort using MPI Send/Recv}


\section*{Message tags}
- Communicating processes may need to send several messages between each other.
- Message tag: differentiate between different messages being sent.


\section*{Message matching}


\section*{Receiving messages}
```

void Comm.Recv(
Object buf,
int offset,
int count,
Datatype datatype,
int src,
int tag
)

```

\section*{A receiver can get a message without knowing:}
- the amount of data in the message,
- the sender of the message,
```

MPI_ANY_SOURCE

```

- or the tag of the message.

\section*{Synchronous Message Passing}

Synchronous send (Ssend)
- waits until complete message can be accepted by the receiving process before completing the send
Synchronous receive (Recv)
- waits until expected message arrives


Synchronous routines can perform two actions
- transfer data
- synchronize processes

\section*{Asynchronous Message Passing}

Send does not wait for actions to complete before returning
- requires local storage for messages
sometimes explicit (programmer needs to care)
sometimes implicit (transparent to the programmer)


In general
- no synchronisation
- allows local progress

\section*{Blocking / nonblocking}

Blocking: return after local actions are complete, though the message transfer may not have been completed
Non-blocking: return immediately
- assumes that data storage to be used for transfer is not modified by subsequent statements until transfer complete


\section*{A nonblocking communication example}


Synchronous / asynchronous vs. blocking / nonblocking
Synchronous / Asynchronous
- about communication between sender and receiver

Blocking / Nonblocking
- about local handling of data to be sent / received

\section*{MPI Send and receive defaults}

\section*{Send}
- blocking,
- synchrony implementation dependent

Danger of Deadlocks.
Don't make any assumptions!
- depends on existence of buffering, performance considerations etc

\section*{Recv}
- blocking

There are a lot of different variations of this in MPI.


\section*{Sources of deadlocks}
- Send a large message from process 0 to process 1
- If there is insufficient storage at the destination, the send must wait for the user to provide the memory space (through a receive)
- What happens with this code?
\begin{tabular}{cl} 
Process 0 & Process 1 \\
\hline Send (1) & \(\operatorname{Send}(0)\) \\
\(\operatorname{Recv}(1)\) & \(\operatorname{Recv}(0)\)
\end{tabular}
- This is called "unsafe" because it depends on the availability of system buffers in which to store the data sent until it can be received

\section*{Some Solutions to the "unsafe" Problem}
- Order the operations more carefully:
\begin{tabular}{cl} 
Process 0 & Process 1 \\
\hline Send (1) & \(\operatorname{Recv}(0)\) \\
\(\operatorname{Recv}(1)\) & \(\operatorname{Send}(0)\)
\end{tabular}
- Supply receive buffer at same time as send:
\begin{tabular}{ll} 
Process 0 & Process 1 \\
\hline Sendrecv (1) & Sendrecv (0)
\end{tabular}

\section*{More Solutions to the "unsafe" Problem}
- Supply own space as buffer for send
\begin{tabular}{ll} 
Process 0 & Process 1 \\
\hline Bsend (1) & Bsend (0) \\
\(\operatorname{Recv}(1)\) & \(\operatorname{Recv}(0)\)
\end{tabular}
- Use non-blocking operations:
\begin{tabular}{ll} 
Process 0 & Process 1 \\
\hline Isend (1) & Isend (0) \\
Irecv (1) & Irecv (0) \\
Waitall & Waitall
\end{tabular}

\section*{MPI is Simple}
- Many parallel programs can be written using just these six functions, only two of which are non-trivial:
- MPI_INIT - initialize the MPI library (must be the first routine called)
- MPI_COMM_SIZE - get the size of a communicator
- MPI_COMM_RANK - get the rank of the calling process in the communicator
- MPI_SEND - send a message to another process
- MPI_RECV - send a message to another process
- MPI_FINALIZE - clean up all MPI state (must be the last MPI function called by a process)
- For performance, however, you need to use other MPI features

\section*{Example: compute Pi}
- The irrational number Pi has many digits
- And it's not clear if they're randomly distributed!
- But they can be computed
\[
\pi \approx h \sum_{i=0}^{N-1} \frac{4}{1+\left(h\left(i+\frac{1}{2}\right)\right)^{2}}
\]
```

for(int i=0; i<numSteps; i++) {
double x=(i + 0.5) * h;
sum += 4.0/(1.0 + x*x);
}
double pi=h * sum ;

```

\section*{Pi record smashed as team finds twoquadrillionth digit}

By Jason Palmer
Science and technology reporter, BBC News
(1) 16 September 2010 | Technology \(f \geqslant \quad \geqslant\) Share

A researcher has calculated the 2,000,000,000,000,000th digit of the mathematical constant pi - and a few digits either side of it.

Nicholas Sze, of tech firm Yahoo, said that when pi is expressed in binary, the two quadrillionth "bit" is 0 .

Mr Sze used Yahoo's Hadoop cloud computing technology to more than
\[
\begin{array}{|l}
\left\{\sum_{0 \leq k<\frac{n+x}{4}} A_{k}+\sum_{\frac{n+x}{4} \leq k} B_{k}\right\} \\
\text { OTHER } \\
\begin{array}{c}
\text { The formula turns an infinite sum into a more } \\
\text { manageable calculation of single terms }
\end{array}
\end{array}
\] double the previous record.

It took 23 days on 1,000 of Yahoo's computers - on a standard PC, the calculation would have taken 500 years.

\section*{Pi's parallel version}
```

MPI.Init(args);
.. // declare and initialize variables (sum=0 etc.)
int size = MPI.COMM_WORLD.Size();
int rank = MPI.COMM_WORLD.Rank();
for(int i=rank; i<numSteps; i=i+size) {
double x=(i + 0.5) * h;
sum += 4.0/(1.0 + x*x);
}
if (rank != 0) {
double [] sendBuf = new double []{sum};
// 1-element array containing sum
MPI.COMM_WORLD.Send(sendBuf, 0, 1, MPI.DOUBLE, 0, 10);
}
else { // rank == 0
double [] recvBuf = new double [1] ;
for (int src=1 ; src<P; src++) {
MPI.COMM_WORLD.Recv(recvBuf, 0, 1, MPI.DOUBLE, src, 10);
sum += recvBuf[0];
}
}
double pi = h * sum; // output pi at rank 0 only!
MPI.Finalize();

```

\section*{COLLECTIVE COMMUNICATION}

Group Communication
Up to here: point-to-point communication

MPI also supports communications among groups of processors - not absolutely necessary for programming (but very nice!)
- but essential for performance

Examples: broadcast, gather, scatter, reduce, barrier, ...

\section*{Collective Computation - Reduce}
```

public void Reduce(java.lang.Object sendbuf,
int sendoffset,
java.lang.Object recvbuf,
int recvoffset,
int count,
Datatype datatype,
Op op,
int root)

```
\begin{tabular}{l|l|l|} 
& & \\
P0 & A & \\
P1 & B & \(\mathrm{A}+\mathrm{B}+\mathrm{C}+\mathrm{D}\) \\
P2 & C & Reduce \\
P3 & D & \\
& & \\
P0 & A & \\
P1 & B & \\
P2 & C & \\
P3 & D & \\
& & A \\
\hline & & \(\mathrm{A}+\mathrm{B}\) \\
\hline
\end{tabular}

\section*{Reduce implementation: a tree-structured global sum}

1. In the first phase:
(a) Process 1 sends to 0,3 sends to 2,5 sends to 4 , and 7 sends to 6 .
(b) Processes 0, 2, 4, and 6 add in the received values.
2. Second phase:
(c) Processes 2 and 6 send their new values to processes 0 and 4, respectively.
(d) Processes 0 and 4 add the received values into their new values.
3. Finally:
(a) Process 4 sends its newest value to process 0 .
(b) Process 0 adds the received value to its newest value.

\section*{Collective Data Movement - Broadcast}


\section*{Collective Computation - Allreduce}
```

public void Allreduce(java.lang.Object sendbuf,
int sendoffset,
java.lang.Object recvbuf,
int recvoffset,
int count,
Datatype datatype,
Op op)

```


Useful in a situation in which all of the processes need the result of a global sum in order to complete some larger computation.

\section*{Allreduce = Reduce + Broadcast?}

Q: What is the number of steps needed?


A global sum followed by distribution of the result.

\section*{Allreduce \(\neq\) Reduce + Broadcast}


Q: What is the number of steps needed?

\section*{Baidu's 'Ring Allreduce' Library \\ Increases Machine Learning Efficiency Across Many GPU Nodes}
by Lucian Armasu February 21, 2017 at 9:10 AM

\section*{© 0 - 0 - ©}

- Baidu's ring alleeduce algorithm

Baidu's silicon Valley AI Lab (SVAIL) announced an implementation of the ring allreduce algorithm for the deep learning community, which will enable significantly faster training of neural networks across GPU models.

\section*{Need For Efficient Parallel Training}

As neural networks have grown to include hundreds of millions or even over a billion parameters the number of GPU nodes needed to do the training has also increased. However, the higher the done by each node. Therefore, the need for algorithms that maximize the performance across the highly parallel system has also increased.

Meet Horovod: Uber's Open
Source Distributed Deep Learning Framework for TensorFlow
By Alex Sergeev and Mike Del Balso

\section*{october 17, 201}


\section*{\begin{tabular}{ll|ll|ll|ll|ll|l}
\hline\(f\) & 1.5 K & \(\boldsymbol{y}\) & 139 & in & 1.3 K & \(Y\) & 6 & 6 & 2 & G* \\
\hline
\end{tabular}}

Over the past few years, advances in deep learning have driven tremendous progress in image processing, speech recognition, and forecasting. At Uber, we apply deep learning across our business; from self-driving research to trip forecasting and fraud prevention, deep learning enables our engineers and data scientists to create better experiences for our users.


Introducing Horovod
The realization that a ring-allreduce approach can improve both usability and serformance motivated us to work on our own implementation to address Uber's ensorFlow needs. We adopted Baidu's draft implementation of the TensorFlow ing-allreduce algorithm and built upon it. We outline our process below
1. We converted the code into a stand-alone Python package called Horovod named after a traditional Russian folk dance in which performers dance with inked arms in a circle, much like how distributed Tensorflow processes use Horovod to communicate with each other. At any point in time, various teams at Uber may be using different releases of TensorFlow. We wanted all teams to be able to leverage the ring-allreduce algorithm without needing to upgrade to the latest version of TensorFlow, apply patches to their versions, or even spend time building out the framework. Having a stand-alone package allowed us to ut the time required to install Horovod from about an hour to a few minutes, depending on the hardware.

We replaced the Baidu ring-allreduce implementation with NCCL. NCCL is NVIDIA's library for collective communication that provides a highly optimized version of ring-allreduce. NCCL 2 introduced the ability to run ring-allreduce across multiple machines, enabling us to take advantage of its many performance boosting optimizations.

\section*{Collective Data Movement - Scatter/Gather}

- Scatter can be used in a function that reads in an entire vector on process 0 but only sends the needed components to each of the other processes.
- Gather collects all of the components of the vector onto destination process, then destination process can process all of the components.

\section*{More Collective Data Movement - some more (16 functions total!)}



\section*{Matrix-Vector-Multiply}

Compute \(\boldsymbol{y}=\boldsymbol{A} \cdot \boldsymbol{x}, \quad\) e.g., \(A=\left[\begin{array}{lll}1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9\end{array}\right] \quad x=\left[\begin{array}{l}10 \\ 20 \\ 30\end{array}\right] \quad \mathrm{y}=\left[\begin{array}{l}A_{1} \cdot x \\ A_{2} \cdot x \\ A_{3} \cdot x\end{array}\right]\)
Assume A and x are available only at rank \(0!\)
Assume \(A\) and \(x\) are available only at rank 0!

\section*{1. Broadcast x}


\section*{Matrix-Vector-Multiply}

Compute \(\boldsymbol{y}=\boldsymbol{A} \cdot \boldsymbol{x}, \quad\) e.g., \(A=\left[\begin{array}{lll}1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9\end{array}\right] \quad x=\left[\begin{array}{c}10 \\ 20 \\ 30\end{array}\right] \quad \mathrm{y}=\left[\begin{array}{l}A_{1} \cdot x \\ A_{2} \cdot x \\ A_{3} \cdot x\end{array}\right]\)
Assume \(A\) and \(x\) are available only at rank 0!

\section*{2. Scatter A}


\section*{Matrix-Vector-Multiply}

Compute \(\boldsymbol{y}=\boldsymbol{A} \cdot \boldsymbol{x}, \quad\) e.g., \(A=\left[\begin{array}{lll}1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9\end{array}\right] \quad x=\left[\begin{array}{l}10 \\ 20 \\ 30\end{array}\right] \quad \mathrm{y}=\left[\begin{array}{l}A_{1} \cdot \cdot x \\ A_{2} \cdot x \\ A_{3} \cdot x\end{array}\right]\)
3. Compute locally
\[
\text { PO } \begin{array}{|l|l|l|l|l|l|}
\hline 1 & 2 & 3 & \hline 10 & 20 & 30 \\
\hline
\end{array}
\]
\[
\begin{array}{ll|l|l|l|l|l|l|}
\hline 4 & 5 & 6 \\
\hline & & 10 & 20 & 30 \\
\hline
\end{array}
\]
\begin{tabular}{ll|l|l|l|l|l|l|l|l|l|l|l|l|l|}
\hline 7 & 8 & 9 & \\
\hline
\end{tabular}

\section*{Matrix-Vector-Multiply}

Compute \(\boldsymbol{y}=\boldsymbol{A} \cdot \boldsymbol{x}, \quad\) e.g., \(A=\left[\begin{array}{lll}1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9\end{array}\right] \quad x=\left[\begin{array}{l}10 \\ 20 \\ 30\end{array}\right] \quad \mathrm{y}=\left[\begin{array}{l}A_{1} \cdot x \\ A_{2} \cdot x \\ A_{3} \cdot x\end{array}\right]\)
4. Gather result y


\section*{Iterations}

Assume we want to apply the matrix-vector product iteratively
\[
y_{n}=A y_{n-1}
\]

Example Application:
Eigenvalue Problem for Probability Matrix, as used in Google's Pagerank algorithm.

Then each process needs the results of other processes after one step.
\(\rightarrow\) Need for Gather + Broadcast in one go.
\(\rightarrow\) If you're clever, you find out how to use reduce_scatter for this ©)!

\section*{Visualizing Program Behavior}


\section*{MPI conclusion}
- The de-facto interface for distributed parallel computing (nearly 100\% market share in HPC)
- Elegant and simple interface
- Definitely simpler than shared memory (no races, limited conflicts, avoid deadlocks with nonblocking communication)
- We only covered the basics here, MPI-3.1 (2015) has 600+ functions
- More concepts:

Derived datatypes
Process topologies
Nonblocking and neighborhood collectives
One-sided accesses (getting the fun of shared memory back ...)
Profiling interfaces

\section*{Sorting}

\section*{(one of the most fun problems in CS)}


\section*{Literature}
- D.E. Knuth. The Art of Computer Programming, Volume 3: Sorting and Searching, Third Edition. Addison-Wesley, 1997. ISBN 0-201-89685-0. Section 5.3.4: Networks for Sorting, pp. 219-247.
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. Introduction to Algorithms, Second Edition. MIT Press and McGraw-Hill, 1990. ISBN 0-262-03293-7. Chapter 27: Sorting Networks, pp.704-724.

\section*{google}
"chapter 27 sorting networks"

\section*{How Fast can we Sort?}

Heapsort \& Mergesort have \(O(n \log n)\) worst-case run time
Quicksort has \(O(n \log n)\) average-case run time
These bounds are all tight, actually \(\Theta(n \log n)\)
So maybe we can dream up another algorithm with a lower asymptotic complexity, such as \(O(n)\) or \(O(n \log \log n)\)
This is unfortunately IMPOSSIBLE!
But why?

\section*{Permutations}

Assume we have n elements to sort
For simplicity, also assume none are equal (i.e., no duplicates)
How many permutations of the elements (possible orderings)?

Example, n=3
\[
\begin{array}{lll}
a[0]<a[1]<a[2] & a[0]<a[2]<a[1] & a[1]<a[0]<a[2] \\
a[1]<a[2]<a[0] & a[2]<a[0]<a[1] & a[2]<a[1]<a[0]
\end{array}
\]

In general, n choices for first, \(\mathrm{n}-1\) for next, \(\mathrm{n}-2\) for next, etc. \(\rightarrow \mathrm{n}(\mathrm{n}-1)(\mathrm{n}-2) \ldots(1)=\mathrm{n}\) ! possible orderings

\section*{Representing Every Comparison Sort}

Algorithm must "find" the right answer among n ! possible answers
Starts "knowing nothing" and gains information with each comparison Intuition is that each comparison can, at best, eliminate half of the remaining possibilities

Can represent this process as a decision tree
- Nodes contain "remaining possibilities"
- Edges are "answers from a comparison"
- This is not a data structure but what our proof uses to represent "the most any algorithm could know"

\section*{Decision Tree for \(\mathbf{n}=\mathbf{3}\)}


The leaves contain all possible orderings of \(a, b, c\)

\section*{What the decision tree tells us}

Binary tree because
- Each comparison has binary outcome
- Assumes algorithm does not ask redundant questions

Because any data is possible, any algorithm needs to ask enough questions to decide among all \(n\) ! answers
- Every answer is a leaf (no more questions to ask)
- So the tree must be big enough to have \(n\) ! leaves
- Running any algorithm on any input will at best correspond to one root-to-leaf path in the decision tree
So no algorithm can have worst-case running time better than the height of the decision tree

\section*{Where are we}

Proven: No comparison sort can have worst-case better than the height of a binary tree with \(n\) ! leaves
- Turns out average-case is same asymptotically
- So how tall is a binary tree with \(n\) ! leaves?

Now: Show a binary tree with \(n\) ! leaves has height \(\Omega(n \log n)\)
- \(n \log n\) is the lower bound, the height must be at least this
- It could be more (in other words, a comparison sorting algorithm could take longer but can not be faster)
Conclude that: (Comparison) Sorting is \(\Omega(\mathrm{n} \log \mathrm{n})\)

\section*{Lower Bound on Height}

The height of a binary tree with \(L\) leaves is at least \(\log _{2} L\)

So the height of our decision tree, \(h\) :

\[
\begin{array}{rlrl}
h & \geq \log _{2}(n!) & & \text { property of binary trees } \\
& =\log _{2}\left(n^{*}(n-1)^{*}(n-2) \ldots(2)(1)\right) & & \text { definition of factorial } \\
& =\log _{2} n+\log _{2}(n-1)+\ldots+\log _{2} 1 & & \text { property of logarithms } \\
& \geq \log _{2} n+\log _{2}(n-1)+\ldots+\log _{2}(n / 2) & & \text { keep first } n / 2 \text { terms } \\
& \geq(n / 2) \log _{2}(n / 2) & & \text { each of the } n / 2 \text { terms left is } \geq \log _{2}(n / 2) \\
& \geq(n / 2)\left(\log _{2} n-\log _{2} 2\right) & & \text { property of logarithms } \\
& \geq(1 / 2) n \log _{2} n-(1 / 2) n & & \text { arithmetic } \\
"=" \Omega\left(n \log ^{n}\right) & &
\end{array}
\]

\section*{Breaking the lower bound on sorting}

Assume 32/64-bit Integer:


Nothing is ever
straightforward in computer
science..
Bogo Sort ( n !)
Stooge Sort ( \(\mathrm{n}^{2.7}\) )
\[
\begin{aligned}
& \text { 3ubble sort } \\
& \text { Shell sort }
\end{aligned}
\]


Radix sort

\section*{SORTING NETWORKS}

\section*{Comparator}
shorter notation:

```

void compare(int[] a, int i, int j, boolean dir) {
if (dir==(a[i]>a[j])){
int t=a[i];
a[i]=a[j];
a[j]=t;
}
}

```


Sorting Networks


Sorting networks are data-oblivious (and redundant)
Data-oblivious comparison tree


\section*{Recursive Construction : Insertion}


\section*{Recursive Construction: Selection}


\section*{Applied recursively..}

insertion sort

bubble sort

with parallelism: insertion sort = bubble sort !

\section*{Question}

How many steps does a computer with infinite number of processors (comparators) require in order to sort using parallel bubble sort? Answer: 2n-3
Can this be improved ?

How many comparisons ?
Answer: (n-1) n/2


How many comparators are required (at a time)?
Answer: n/2
Reusable comparators: \(\mathrm{n}-1\)

\section*{Improving parallel Bubble Sort}

Odd-Even Transposition Sort:
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline 0 & 9 & \(\leftrightarrow\) & 8 & & 2 & \(\leftrightarrow\) & 7 & & 3 & \(\leftrightarrow\) & 1 & & 5 & \(\leftrightarrow\) & 6 & \\
\hline 1 & 8 & & 9 & \(\leftrightarrow\) & 2 & & 7 & \(\leftrightarrow\) & 1 & & 3 & \(\leftrightarrow\) & 5 & & 6 & \(\leftrightarrow\) \\
\hline 2 & 8 & \(\leftrightarrow\) & 2 & & 9 & \(\leftrightarrow\) & 1 & & 7 & \(\leftrightarrow\) & 3 & & 5 & \(\leftrightarrow\) & 4 & \\
\hline 3 & 2 & & 8 & \(\leftrightarrow\) & 1 & & 9 & \(\leftrightarrow\) & 3 & & 7 & \(\leftrightarrow\) & 4 & & 5 & \(\leftrightarrow\) \\
\hline 4 & 2 & \(\leftrightarrow\) & 1 & & 8 & \(\leftrightarrow\) & 3 & & 9 & \(\leftrightarrow\) & 4 & & 7 & \(\leftrightarrow\) & 5 & \\
\hline 5 & 1 & & 2 & \(\leftrightarrow\) & 3 & & 8 & \(\leftrightarrow\) & 4 & & 9 & \(\leftrightarrow\) & 5 & & 7 & \(\leftrightarrow\) \\
\hline 6 & 1 & \(\leftrightarrow\) & 2 & & 3 & \(\leftrightarrow\) & 4 & & 8 & \(\leftrightarrow\) & 5 & & 9 & \(\leftrightarrow\) & 6 & \\
\hline 7 & 1 & & 2 & \(\leftrightarrow\) & 3 & & 4 & \(\leftrightarrow\) & 5 & & 8 & \(\leftrightarrow\) & 6 & & 9 & \(\leftrightarrow\) \\
\hline 8 & 1 & \(\leftrightarrow\) & 2 & & 3 & \(\leftrightarrow\) & 4 & & 5 & \(\leftrightarrow\) & 6 & & 8 & \(\leftrightarrow\) & 7 & \\
\hline & 1 & & 2 & & 3 & & 4 & & 5 & & 6 & & 7 & & 8 & \\
\hline
\end{tabular}
void oddEvenTranspositionSort(int[] a, boolean dir) \{ int \(n=\) a.length;
    for (int \(i=0 ; i<n ;++i\) ) \{
        for (int \(j=1 \% 2 ; j+1<n ; j+=2)\)
            compare(a,j,j+1,dir);
    \}
\}


\section*{Improvement?}


Same number of comparators (at a time)
Same number of comparisons
But less parallel steps (depth): n

In a massively parallel setup, bubble sort is thus not too bad.

But it can go better..

\section*{How to get to a sorting network?}
- It's complicated ()
- In fact, some structures are clear but there is a lot still to be discovered!

\section*{- For example:}
- what is the minimum number of comparators?
- What is the minimum size?
- Tradeoffs between these two?

Optimal sorting networks
[edit]

Source: wikipedia
For small, fixed numbers of inputs \(n\), optimal sorting networks can be constructed, with either minimal depth (for maximally parallel execution) or minimal size (number of comparators). These networks can be used to increase the performance of larger sorting networks resulting from the recursive constructions of, e.g., Batcher, by halting the recursion early and inserting optimal nets as base cases. \({ }^{[9]}\) The following table summarizes the known optimality results:
\begin{tabular}{|l|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline \(\boldsymbol{n}\) & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 & 17 \\
\hline Depth \(^{[10]}\) & 0 & 1 & 3 & 3 & 5 & 5 & 6 & 6 & 7 & 7 & 8 & 8 & 9 & 9 & 9 & 9 & 10 \\
\hline Size, upper bound \(^{[11]}\) & 0 & 1 & 3 & 5 & 9 & 12 & 16 & 19 & 25 & 29 & 35 & 39 & 45 & 51 & 56 & 60 & 71 \\
\hline\({\text { Size, lower bound (if different) }{ }^{[11]}}^{[11]}\) & & & & & & & & & & 33 & 37 & 41 & 45 & 49 & 53 & 58 \\
\hline
\end{tabular}

The first sixteen depth-optimal networks are listed in Knuth's Art of Computer Programming, \({ }^{[1]}\) and have been since the 1973 edition; however, while the optimality of the first eight was established by Floyd and Knuth in the 1960s, this property wasn't proven for the final six until 2014 \({ }^{[12]}\) (the cases nine and ten having been decided in \(1991{ }^{[9]}\) ).
For one to ten inputs, minimal (i.e. size-optimal) sorting networks are known, and for higher values, lower bounds on their sizes \(S(n)\) can be derived inductively using a lemma due to Van Voorhis: \(S(n+1) \geq S(n)+\left\lceil\log _{2}(n)\right\rceil\). All ten optimal networks have been known since 1969, with the first eiaht again being known as optimal since the work of Floyd and Knuth, but optimality of the cases \(n=9\) ard \(n=10\) took until 2014 to eresolved. \({ }^{[11]}\)

\section*{Interlude: Machine Models}

RAM : Random Access Machine
- Unbounded local memory
- Each memory has unbounded capacity
- Simple operations: data, comparison, branches
- All operations take unit time

Time complexity: number of steps executed

Space complexity: (maximum) number of memory cells used

\section*{Machine Models}

PRAM : Parallel Random Access Machine
- Abstract machine for designing algorithms applicable for parallel computers
- Unbounded collection of RAM processors \(P_{0}, P_{1}, \ldots\)
- Each processor has unbounded registers

- Unbounded shared memory
- All processors can access all memory in unit time
- All communication via shared memory
shared memory

\section*{Shared Memory Access Model}

ER: processors can simultaneously read from distinct memory locations EW: processors can simultaneously write to distinct memory locations CR: processors can simultanously read from any memory location CW: processors can simultaneously write to any memory location Specification of the machine model as one of EREW, CREW, CRCW

\section*{Example: Why the machine model can be important}

Find maximum of \(\mathbf{n}\) elements in an array \(\mathbf{A}\)
Assume \(O\left(n^{2}\right)\) processors and the CRCW model
For all \(i \in\{0,1, \ldots, n-1\}\) in parallel do
\[
P_{i 0}: m_{i} \leftarrow \text { true }
\]

For all \(i, j \in\{0,1, \ldots, n-1\}, i \neq j\) in parallel do
\[
P_{i j}: \text { if } A_{i}<A_{j} \text { then } m_{i} \leftarrow \text { false }
\]

For all \(i \in\{0,1, \ldots, n-1\}\) in parallel do
\[
P_{i 0}: \text { if } m_{i}=\text { true then } \max \leftarrow A_{i}
\]
\(\mathrm{O}(1)\) time complexity!

Illustration
1. Init


\section*{CREW}

Q: How many steps does max-find require with CREW?
Using CREW only two values can be merged into a single value by one processor at a time step: number of values that need to be merged can be halved at each step \(\rightarrow\) Requires \(\Omega(\log n)\) steps

There are a lot of interesting theoretical results for PRAM machine models (e.g., CRCW simulatable with EREW) and for PRAM based algorithms (e.g., cost optimality / time optimality proofs etc). We will not go into more details here.

In the following we assume a CREW PRAM model -- and receive in retrospect a justification for the results stated above on parallel bubble sorting.

\section*{Parallel sorting}


Prove that the two networks above sort four numbers. Easy?

\section*{Zero-one-principle}

Theorem: If a network with \(n\) input lines sorts all \(2^{n}\) sequences of 0 s and 1 s into non-decreasing order, it will sort any arbitrary sequence of \(n\) numbers in nondecreasing order.

\section*{Proof}
\begin{tabular}{|l|l|l|l|l|l|}
\hline 1 & 8 & 20 & 30 & 5 & 9 \\
\hline
\end{tabular}\(\Rightarrow\)\begin{tabular}{|l|l|l|l|l|l|}
\hline 1 & 5 & 8 & 9 & 20 & 30 \\
\hline
\end{tabular}

Argue: If \(x\) is sorted by a network \(N\) then also any monotonic function of \(x\).
e.g., floor \((x / 2)\)\begin{tabular}{|l|l|l|l|l|l|}
\hline 0 & 4 & 10 & 15 & 2 & 4 \\
\cline { 2 - 10 }
\end{tabular}\(\quad\)\begin{tabular}{|l|l|l|l|l|l|}
\hline 0 & 2 & 4 & 4 & 10 & 15 \\
\hline
\end{tabular}

Show: If x is not sorted by the network, then there is a monotonic function \(f\) that maps \(x\) to 0 s and 1 s and \(f(x)\) is not sorted by the network

\begin{tabular}{|l|l|l|l|l|l|}
\hline 0 & 0 & 1 & 1 & 0 & 1 \\
\hline
\end{tabular}\(\Rightarrow\)\begin{tabular}{|l|l|l|l|l|l|}
\hline 0 & 0 & 1 & 0 & 1 & 1 \\
\hline
\end{tabular}

\section*{\(x\) not sorted by \(N \Rightarrow\) there is an \(f(x) \in\{0,1\}^{n}\) not sorted by N \(\Leftrightarrow\)}
\(f\) sorted by N for all \(f \in\{0,1\}^{n} \Rightarrow x\) sorted by N for all x

\section*{Proof}

Assume a monotonic function \(f(x)\) with \(f(x) \leq f(y)\) whenever \(x \leq y\) and a network \(N\) that sorts. Let \(\mathbf{N}\) transform \(\left(x_{1}, x_{2}, \ldots, x_{n}\right)\) into \(\left(y_{1}, y_{2}, \ldots, y_{n}\right)\), then it also transforms \(\left(f\left(x_{1}\right), f\left(x_{2}\right), \ldots, f\left(x_{n}\right)\right)\) into \(\left(f\left(y_{1}\right), f\left(y_{2}\right), \ldots, f\left(y_{n}\right)\right)\).

All comparators must act in the same way for the \(f\left(x_{i}\right)\) as they do for the \(x_{i}\)
Assume \(y_{i}>y_{i+1}\) for some \(i\), then consider the monotonic function
\[
f(x)=\left\{\begin{array}{l}
0, \text { if } x<y_{i} \\
1, \text { if } x \geq y_{i}
\end{array}\right.
\]
\(\rightarrow \mathbf{N}\) converts
\(\left(f\left(x_{1}\right), f\left(x_{2}\right), \ldots, f\left(x_{n}\right)\right)\) into \(\left(f\left(y_{1}\right), f\left(y_{2}\right), \ldots f\left(y_{i}\right), f\left(y_{i+1}\right), \ldots, f\left(y_{n}\right)\right)\)

\section*{Bitonic Sort}

Bitonic (Merge) Sort is a parallel algorithm for sorting
If enough processors are available, bitonic sort breaks the lower bound on sorting for comparison sort algorithm
Time complexity of \(O\left(n \log ^{2} n\right)\) (sequential execution)
Time complexity of \(O\left(\log ^{2} n\right)\) (parallel time)
Worst \(=\) Average \(=\) Best case

\section*{What is a Bitonic Sequence?}


Monotonic ascending sequence


Monotonic descending sequence

\section*{Bitonic Sets}


A bitonic set is defined as a set in which the sign of the gradient changes once at most.
So that \(x_{0} \leq \cdots \leq x_{k} \geq \cdots \geq x_{n-1}\), for some \(k, 0 \leq k<n\)

\section*{Bitonic sets - Wraparound}

(a) Single maximum

(b) Single maximum and single minimum

A bitonic sequence is defined as a list with no more than one Local maximum and no more than one Local minimum.

\section*{Bitonic (again)}

Sequence ( \(x_{1}, x_{2}, \ldots, x_{n}\) ) is bitonic, if it can be circularly shifted such that it is first monotonically increasing and then monontonically decreasing.
\[
(1,2,3,4,5,3,1,0) \quad(4,3,2,1,2,4,6,5)
\]

Bitonic 0-1 Sequences

\section*{\(0^{i} 1^{j} 0^{k}\)}
\(\ldots 1^{i} 0^{j} 1^{k}\)

\section*{Properties}

If ( \(x_{1}, x_{2}, \ldots, x_{n}\) ) is monotonically increasing (decreasing) and then monotonically decreasing (increasing), then it is bitonic If \(\left(x_{1}, x_{2}, \ldots, x_{n}\right)\) is bitonic, then \(\left(x_{1}, x_{2}, \ldots, x_{n}\right)^{R}:=\left(x_{n}, x_{n-1}, \ldots, x_{1}\right)\) is also bitonic

\section*{The Half-Cleaner}


\section*{The Half-Cleaner}

void halfClean(int[] a, int lo, int m, boolean dir)
\{
for (int \(i=10 ; i<l o+m ; i++\) ) compare(a, i, i+m, dir);
\}


\section*{Binary Split: Application of the Half-Cleaner}
1. Divide the bitonic list into two equal halves.
2. Compare-Exchange each item on the first half with the corresponding item in the second half.


\section*{Binary splits - Result}

Two bitonic sequences where the numbers in one sequence are all less than the numbers in the other sequence.
Because the original sequence was bitonic, every element in the lower half of new sequence is less than or equal to the elements in its upper half.

\section*{Sequence D}


\section*{Bitonic Split Example}


\section*{Lemma}

Input bitonic sequence of 0 s and 1 s , then for the output of the half-cleaner it holds that
- Upper and lower half is bitonic
- One of the two halfs is bitonic clean
- Every number in upper half \(\leq\) every number in the lower half

\section*{Proof: All cases}





The four remaining cases (010 \(\rightarrow\) 101)


\section*{Construction of a Bitonic Sorting Network}
bitonic


\section*{Recursive Construction}

```

void bitonicMerge(int[] a, int lo, int n, boolean dir)
{
if (n>1){
int m=n/2;
halfClean(a, lo, m, dir);
bitonicMerge(a, lo, m, dir);
bitonicMerge(a, lo+m, m, dir);
}
}

```


\section*{Bitonic Merge}
- Compare-and-exchange moves smaller numbers of each pair to left and larger numbers of pair to right.
- Given a bitonic sequence, recursively performing 'binary split' will sort the list.

Bitonic sequence


\section*{Bi-Merger}


Bi-Merger on two sorted sequences acts like a half-cleaner on a bitonic sequence (when one of the sequences is reversed)

Merger


\section*{Recursive Construction of a Sorter}


GHzürich
```

private void bitonicSort(int a[], int lo, int n, boolean dir) {
if (n>1){
int m=n/2;
bitonicSort(a, lo, m, ASCENDING);
bitonicSort(a, lo+m, n, DESCENDING);
bitonicMerge(a, lo, n, dir);
}
}

```


\section*{Example}


\section*{Example}


\section*{Bitonic Merge Sort}

How many steps?
```

\#mergers

```
\[
\sum_{i=1}^{\log n} \log 2^{i}=\sum_{i=1}^{\log n} i \log 2=\frac{\log n \cdot(\log n+1)}{2}=O\left(\log ^{2} n\right)
\]```

