Parallel Computing

A brief overview

Disclaimer: *These abbreviated notes DO NOT substitute the textbook for this class. They should be used IN CONJUNCTION with the textbook and the material presented in class. If there is a discrepancy between these notes and the textbook, ALWAYS consider the textbook to be correct. Report such a discrepancy to the instructor so that he resolves it. These notes are only distributed to the students taking this class with A. Gerbessiotis in Spring 2012 ; distribution outside this group of students is NOT allowed.*

(c) Copyright A. Gerbessiotis. All rights reserved.

1. What is a parallel algorithm?

A *parallel algorithm* is an algorithm designed for a parallel computer.

2. What is a parallel computer?

A *parallel computer* is a collection of processors that cooperatively solve computationally intensive problems faster than other computers. Parallel algorithms allow the efficient programming of parallel computers. This way the waste of computational resources can be avoided.

The term supercomputer refers to a general-purpose computer that can solve computational intensive problems faster than traditional computers. A supercomputer may or may not be a parallel computer.

The cost of building a parallel computer has dropped significantly in the past years. One is able to build a parallel computer using off-the-shelf commercial components. This way workstations can be arranged together in a cluster/network and interconnected with a high-speed network (eg Myrinet siwtches, 100Mb/sec or Gigabit ethernet switches) and form a high-performance parallel machine. The construction of a parallel computer is more complicated than that of a sequential machine. One needs to take into consideration not only those issues related with building a sequential machine but also issues related with the interconnection of multiple processors and their efficient coordination. A complete network where each processor can directly communicate with every other processors is not currently technologically feasible if the parallel machine consists of more than a few processors. Methods thus need to be devised that allow the efficient dispatch of information from one processor to another (routing protocols and algorithms). Problems like congestion (bottlenecks) may occur if someone is not very careful in designing such methods.

(c) Copyright A. Gerbessiotis. All rights reserved.

Flynn's Taxonomy of computer architectures (control mechanism)
Depending on the execution and data streams computer architectures can be distinguished into the following groups.

(1) *SISD (Single Instruction Single Data)* : This is a sequential computer.

(2) *SIMD (Single Instruction Multiple Data)* : This is a parallel machine like the TM CM-200. SIMD machines are suited for **data-parallel** programs programs where the same set of instructions are executed on a large data set. More recently GPU (Graphics Processor Units) such as those supporting NVIDIA CUDA (Compute Unified Device Architecture) fall into this model. Intel MMX instruction set supports SIMD execution.

(3) *MISD (Multiple Instructions Single Data)* : Some consider a systolic array a member of this group. Or the pipelines of modern microprocessors.

(4) *MIMD (Multiple Instructions Multiple Data)* : All other parallel machines. A MIMD architecture can be an MPMD or an SPMD. In a Multiple Program Multiple Data organization, each processor executes its own program as opposed to a single program that is executed by all processors on a Single Program Multiple Data architecture.

Some consider CM-5 as a combination of a MIMD and SIMD as it contains control hardware that allows it to operate in a SIMD mode.

(c) Copyright A. Gerbessiotis. All rights reserved.

Performance Characteristics of a Parallel Algorithm
The performance of a parallel algorithm A that solves some problem instance can be measured in terms of various measures.

- *Parallel Time* T gives the execution time of the parallel algorithm and $T = max_i T_i$, where T_i is the running time of algorithm A on processor i.
- *Processor size* P is the number of processors assigned to solve a particular problem instance.
- *Work* is the product $P \cdot T$. Sometimes the actual work $\sum_{j=1}^{T} P_j \leq W$ is used instead, where P_j is the number of processors that are active in step j.
- *Speedup* given by $s = T_s/T$ is the ratio of the execution time T_s of the most efficient sequential algorithm that solves a particular problem instance to parallel time T of A on the same instance.
- *Scaled speedup* is the ratio T_1/T_p, where T_i here denotes the parallel time T of A when i processors are used to solve that particular problem instance.
- *Efficiency* $e = s/p$, which is sometimes expressed as a percentage. It is also the case that $e = T_s/W$, where $W = TP$.
- *Scaled Efficiency*, which is derived similarly from scaled speedup.

(c) Copyright A. Gerbessiotis. All rights reserved.

The PRAM Model

The Parallel Random Access Machine (PRAM) is one of the simplest ways to model a parallel computer. A PRAM consists of a collection of (sequential) processors that can *synchronously* access a global *shared* memory in unit time. Each processor can thus access its shared memory as fast (and efficiently) as it can access its own local memory. The main advantages of the PRAM is its simplicity in capturing parallelism and abstracting away communication and synchronization issues related to parallel computing. Processors are considered to be in abundance and unlimited in number. The resulting PRAM algorithms thus exhibit *unlimited parallelism* (number of processors used is a function of problem size). The abstraction thus offered by the PRAM is a fully synchronous collection of processors and a shared memory which makes it popular for parallel algorithm design. It is, however, this abstraction that also makes the PRAM unrealistic from a practical point of view. Full synchronization offered by the PRAM is too expensive and time demanding in parallel machines currently in use. Remote memory (i.e. shared memory) access is considerably more expensive in real machines than local memory access as well and UMA machines with unlimited parallelism are difficult to build.

Depending on how concurrent access to a single memory cell (of the shared memory) is resolved, there are various PRAM variants. ER (Exclusive Read) or EW (Exclusive Write) PRAMs do not allow concurrent access of the shared memory. It is allowed, however, for CR (Concurrent Read) or CW (Concurrent Write) PRAMs. Combining the rules for read and write access there are four PRAM variants: EREW, ERCW, CREW and CRCW PRAMs. Moreover, for CW PRAMs there are various rules that arbitrate how concurrent writes are handled.

Convention: In this subject we name processors arbitrarily either $0, 1, \ldots, p-1$ or $1, 2, \ldots, p$.

(c) Copyright A. Gerbessiotis. All rights reserved.

(1) in the *arbitrary* PRAM, if multiple processors write into a single shared memory cell, then an arbitrary processor succeeds in writing into this cell,

(2) in the *common* PRAM, processors must write the same value into the shared memory cell,

(3) in the *priority* PRAM the processor with the highest priority (smallest or largest indexed processor) succeeds in writing,

(4) in the *combining* PRAM if more than one processors write into the same memory cell, the result written into it depends on the combining operator. If it is the *sum* operator, the sum of the values is written, if it is the *maximum* operator the maximum is written.

The EREW PRAM is the weakest among the four basic variants. A CREW PRAM can simulate an EREW one. Both can be simulated by the more powerful CRCW PRAM. An algorithm designed for the common PRAM can be executed on a priority or arbitrary PRAM and exhibit similar complexity. The same holds for an arbitrary PRAM algorithm when run on a priority PRAM.

Assumptions In this handout we examine parallel algorithms on the PRAM. In the course of the presentation of the various algorithms some common assumptions will be made. The input to a particular problem would reside in the cells of the shared memory. We assume, in order to simplify the exposition of our algorithms, that a cell is wide enough (in bits or bytes) to accommodate a single instance of the input (eg. a key or a floating point number). If the input is of size n, the first n cells numbered $0, \ldots, n-1$ store the input. In the discussion below, we assume that the number of processors of the PRAM is n or a polynomial function of the size n of the input. Processor indices are $0, 1, \ldots, n-1$.

(c) Copyright A. Gerbessiotis. All rights reserved.

Parallel Sum operation

What happens if we have a parallel machine? How fast can he evaluate a polynomial? Can we evaluate a polynomial $f(x)$ p times faster than the sequential time? If we have more than one processors, we can thus assume that the sequential model of computation that we have been using can be extended so that all p processors have access to a shared (common) memory, that they can read in unit time.

Problem: Parallel Sum.

Input. $x_0, \ldots, x_{n-1}$

Output. Evaluate $x_0 + \ldots + x_{n-1}$.

A sequential algorithm that solves this problem requires $n-1$ additions. For a PRAM implementation, value x_i is initially stored in shared memory location $M[i]$. The sum $x_0 + x_1 + \ldots + x_{n-1}$ is to be computed in $T = \lg n$ parallel steps. Without loss of generality, we assume that n is a power of two. The algorithm consists of $\lg n$ steps. In step i, processor $j < n/2^i$ reads shared-memory locations $M[2j]$ and $M[2j+1]$ and sums these values and stores the result into memory cell j. After $\lg n$ steps the sum resides in cell 0. Algorithm PARALLEL_SUM below has $T = O(\lg n)$, $P = n/2$, and $W = O(n \lg n)$, $W_2 = O(n)$.

```
//   pid() returns the id of the processor issuing the call.
begin PARALLEL_SUM  (n)
1.    i = 1 ; j = pid();
2.    while (j < n/2^i)
3.        a = M[2j];
4.        b = M[2j + 1];
5.        M[j] = a + b;
6.        i = i + 1;
7.    end
end PARALLEL_SUM
```

(c) Copyright A. Gerbessiotis. All rights reserved.

An Example

Algorithm Parallel Sum.

```
M[0]      M[1]      M[2]      M[3]      M[4]      M[5]      M[6]      M[7]

 x0        x1        x2        x3        x4        x5        x6        x7     t=0

x0+x1     x2+x3     x4+x5     x6+x7                                            t=1

x0+...+x3 x4+...+x7                                                          t=2

x0+...+x7                                                                    t=3
```

Algorithm PARALLEL_SUM can be easily extended to include the case where n is not a power of two. PARALLEL_SUM is the first instance of a sequential problem that has a trivial sequential but more complex parallel solution. Instead of operator *Sum* other operators like *Multiply*, *Maximum*, *Minimum*, or in general, any associative operator could have been used. As associative operator $\otimes$ is one such that $(a \otimes b) \otimes c = a \otimes (b \otimes c)$.

Exercise 1 *Can you improve* PARALLEL_SUM *so that T remains the same, $P = O(n/\lg n)$, and $W = O(n)$? Explain.*

Exercise 2 *What if i have p processors where $p < n$? (You may assume that n is a multiple of p).*

Exercise 3 *Generalize the Parallel Sum algorithm to any associative operator.*

(c) Copyright A. Gerbessiotis. All rights reserved.

Broadcasting

A word is stored in memory location M[0] of the shared memory. We would like this word to be read by all n processors of a PRAM. On a CREW PRAM this requires one parallel step (processor i concurrently reads cell 0). On an EREW PRAM broadcasting can be performed in $O(\lg n)$ steps. The structure of the algorithm is the reverse of the previous one. In $\lg n$ steps the word is broadcast as follows. In step i each processor with index j less than 2^i reads the contents of cell $M[j]$ and copies it into cell $M[j + 2^i]$. After $\lg n$ steps each processor i reads the message by reading the contents of cell i. The EREW PRAM algorithm below that solves the broadcasting problem has performance $P = O(n)$, $T = O(\lg n)$, and $W = O(n \lg n)$, $W_2 = O(n)$.

```
begin BROADCAST (M)
1.    i = 0 ; j = pid(); M[0]=M;
2.    while (2^i < P)
3.        if (j < 2^i)
5.            M[j + 2^i] = M[j];
6.        i = i + 1;
6.    end
7.    Processor j reads M from M[j].
end BROADCAST
```

Parallel Prefix

Given a set of n values $x_0, x_1, \ldots, x_{n-1}$ and an associative operator, say $\oplus$, the *parallel prefix* problem is to compute the following n results.

0: x_0,

1: $x_0 \oplus x_1$,

2: $x_0 \oplus x_1 \oplus x_2$,

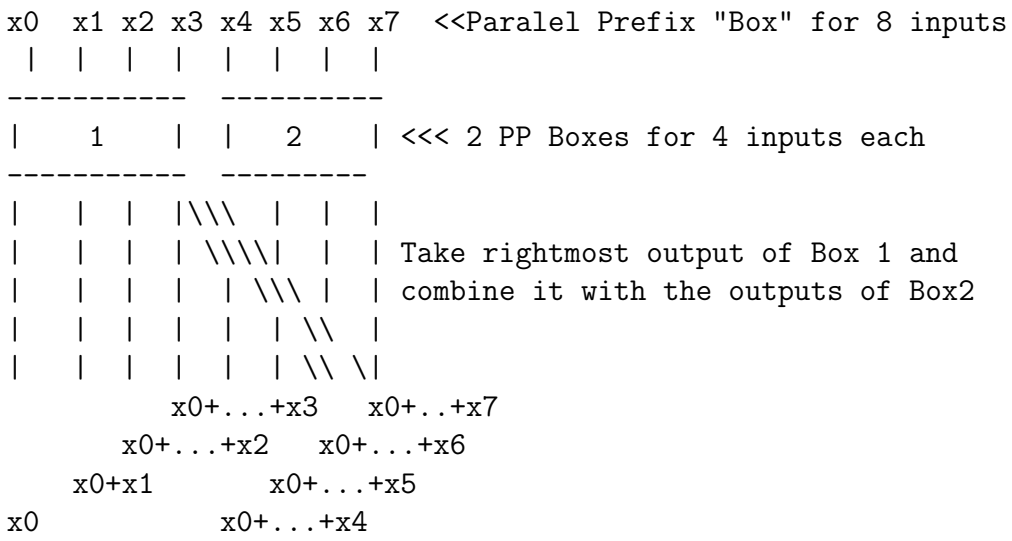
...

$n-1$: $x_0 \oplus x_1 \oplus \ldots \oplus x_{n-1}$.

Parallel prefix is also called *prefix sums* or *scan*. It has many uses in parallel computing such as in load-balancing the work assigned to processors and compacting data structures such as arrays. We shall prove that computing ALL THE SUMS is no more difficult that computing the single sum $x_1 \oplus \ldots \oplus x_n$. For the examples of the remainder, $\oplus$ will be our $+$.

(c) Copyright A. Gerbessiotis. CS 667-101 : Spring 2012 . All rights reserved.

A divide and Conquer Approach: Concurrency required



An algorithm for parallel prefix on a ?REW PRAM would require $\lg n$ phases. In phase i, processor j reads the contents of cells $M[j]$ and $M[j - 2^i]$ (if this cell exists) combines them and stores the result in cell $M[j]$.

The algorithm utilizes the bracketed notation of the following

(c) Copyright A. Gerbessiotis. All rights reserved.

Another algorithm: Concurrency avoided

```
// We write below[1:2] to denote X[1]+X[2]
//               [i:j] to denote X[i]+X[i+1]+...+X[j]
//               [i:i] is X[i]  NOT X[i]+X[i]
//               [1:2][3:4]=[1:2]+[3:4]= (X[1]+X[2])+(X[3]+X[4])=X[1]+X[2]+X[3]+X[4]
// Input :  M[j]= X[j]=[j:j]             for j=1,...,n.
// Output:  M[j]= X[1]+...+X[j] = [1:j]  for j=1,...,n.
 ParallelPrefix(n) // .
1.i=1;             // At this step  M[j]= [j:j]=[j+1-2**(i-1):j]
2.while (i < n ) {
3. j=pid();
4. if (j-2**(i-1) >0 ) {
5.  a=M[j];         // Before this stepM[j]         = [j+1-2**(i-1):j]
6.  b=M[j-2**(i-1)];// Before this stepM[j-2**(i-1)]= [j-2**(i-1)+1-2**(i-1):j-2**(i-1)]
7.  M[j]=a+b;       // After this step M[j]= M[j]+M[j-2**(i-1)]=[j-2**(i-1)+1-2**(i-1):j-2**(i-1)]
                    //                                         [j+1-2**(i-1):j] =
                    //                                         [j-2**(i-1)+1-2**(i-1):j]=
                    //                                         [j+1-2**i:j]
8. }
9. i=i*2;
  }
```

At step 5, memory location $j - 2^{i-1}$ is read provided that $j - 2^{i-1} \geq 1$. This is true for all times $i \leq t_j = \lg(j-1) + 1$. For $i > t_j$ the test of line 4 fails and lines 5-8 are not executed.

Example

For visualization purposes, the second step is written in two different lines. When we write $x_1 + \ldots + x_5$ we mean $x_1 + x_2 + x_3 + x_4 + x_5$.

```
   x1       x2         x3         x4          x5         x6       x7        x8


1.         x1+x2       x2+x3     x3+x4       x4+x5     x5+x6     x6+x7      x7+x8

2.                 x1+(x2+x3)           (x2+x3)+(x4+x5)    (x4+x5)+(x6+x7)
2.                           (x1+x2)+(x3+x4)     (x3+x4)+(x5+x6)    (x5+x6+x7+x8)

3.                                       x1+...+x5          x1+...+x7
3.                                                 x1+...+x6           x1+...+x8
Finally
F. x1   x1+x2    x1+...+x3  x1+...+x4  x1+...+x5 x1+...+x6 x1+...+x7 x1+...+x8
```

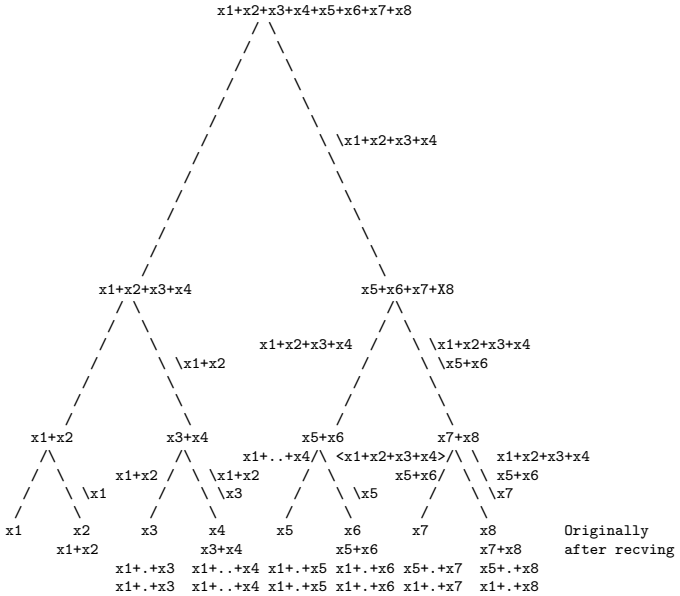
(c) Copyright A. Gerbessiotis. All rights reserved.

Parallel Prefix : Yet another algorithm

Consider the following variation of parallel prefix on n inputs that works on a complete binary tree with n leaves (assume n is a power of two).

```
Action by nodes
 1. Non-leaf  : If it receives l and r from left and right children, computes l + r and
                sends it up and send down to its right child the l.
 2. Root      : Step [1] except nothing is sent up.
 3. Non-leaf  : If it gets p from parent it transmits it to its left/right children.
 4. Leaf      : If it holds l and receives p from its parent it sets l = p + l (this order)
                [note p is the left argument, l is the right one, order matters]
```

```
\begin{mmytinylisting}
\begin{verbatim}
```



Parallel Prefix : A recursive version

The parallel prefix algorithm of the previous page (tree-based) requires about $2\lg n+1$ parallel steps, $P = n$ processors and work $W = \Theta(n \lg n)$, and $W_2 = \Theta(n)$. One could describe that version due to Ladner and Fischer as follows. By rescheduling the computation and using $P = n/\lg n$ processors, the work can be reduced to linear.

```
begin PPF_RECURSIVE  (In[0..n-1],Out[0..n-1],p = 0..n-1)
1.    Out[0] = In[0];
2.    if  n > 1       then
3.        ∀ i = 0,...,n-1      dopar
4.            X[i] = In[2i] + In[2i+1];
5.        enddo
6.        Y=PPF_recursive(X[0..n/2-1],Y[0..n/2-1],p = 0..n/2-1);
7.        ∀ i = 0,...,n/2-1   dopar
8.            Out[2i+1]=Y[i];
9.        enddo
10.       ∀ i = 1,...,n/2-1   dopar
11.           Out[2i]=Y[i-1]+X[2i];
12.       enddo
13.   endif
end PPF_RECURSIVE
```

(c) Copyright A. Gerbessiotis. All rights reserved.

Parallel Prefix : An iterative version

An iterative version of that algorithm is depicted below.

```
begin PPF_ITERATIVE (In[0..n-1],Out[0..n-1],p = 0..n-1)
1.   for i = 0,...,n-1 dopar
2.        T[0,i] = In[i];
3.   enddo
4.   for j = 1,...,lg n   do
5.        for i = 0,...,n/2^j - 1        dopar
6.             T[j,i] = T[j-1,2i] + T[j-1,2i+1];
7.   enddo
8.   for j = lg n,...,0          do
9.        for i = 0dopar
10.            V[j,0] = T[j,0]; //Processor 0 executes only
11.       for odd(i), 0 ≤ i ≤ n/2^j - 1   dopar
12.            V[j,i] = V[j+1,i/2]; //Processor odd(i) executes only
11.       for even(i), 2 ≤ i ≤ n/2^j - 1 dopar
12.            V[j,i] = V[j+1,(i-1)/2]+T[j,i]; //Processor even(i) executes only
13.  enddo
14.  Out[i]=V[0,i];
end PPF_ITERATIVE
```

(c) Copyright A. Gerbessiotis. All rights reserved.

- Add two n-bit binary numbers in $2\lg n + 1$ steps using an n-leaf complete binary tree. The structure on page 13 could be viewed as an inverted binary tree (root at the bottom). However the structure of the computations do not correspond to a complete binary tree. If in the structure of page 13 parallel prefix requires T time, then in a complete binary tree requires 2T time.

Sequential algorithm requires n steps.

	16	15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00
a	0	1	0	1	1	1	0	0	1	0	0	1	0	0	1	0	
b	0	1	1	0	1	0	0	0	0	1	0	1	1	1	0	0	
	s	g	p	p	g	p	s	s	p	p	s	g	p	p	p	s	s
c																	
$a+b$																	

Table 1: Binary addition example

$(a+b)_i = a_i \oplus b_i \oplus c_{i-1}$, where $\oplus = XOR$.

- s: stops a carry bit $(0+0)$
- g: generates a carry bit $(1+1)$
- p: propagates a carry bit $(0+1$ or $1+0)$.

Problem: In order to compute the k-th bit the $k-1$-st carry needs to be computed as well. There exists a non-trivial non-obvious parallel solution.

We shall try for each bit position to find the carry bit required for addition so that all bit positions can be added in parallel. We shall show that carry computation takes $\Theta(\lg n)$ time on a binary tree with a computation that is known to us: parallel prefix.

Observation. The i-th carry bit is one if the leftmost non-p to the right of the i-th bit is a g.

Question. How can we find i-th carry bit?

The previous observation takes the following algorithmic form (figure on the left). Such a computation requires $O(n)$ time for $j = n$ (n-th bit). Let the i-th bit position symbol (p, s, g) be denoted by x_i. Then

```
Scan for j=i, ... , 0
  if p ignore else
    if g
       carry=1
       exit;
    else
       carry=0
       exit;
```

$$c_0 = x_0 = s$$
$$c_1 = x_0 \otimes x_1$$
$$c_2 = x_0 \otimes x_1 \otimes x_2$$
$$c_{16} = x_0 \otimes \ldots \otimes x_{16},$$

where

$\otimes$	s	p	g
s	s	s	g
p	s	p	g
g	s	g	g

Algorithm for parallel addition

Step 1. Compute symbol ($\{s, p, g\}$) for i bit in parallel for all i.

Step 2. Perform a parallel prefix computation on the n symbols plus 0-th symbol s in parallel where operator is defined as in previous table.

Step 3. Combine (exclusive OR) the carry bit from bit position $i-1$ (interpret g as an 1 and an s as a 0) with the exclusive OR of bits in position i to find the i-th bit of the sum.

Steps 1 and 3 require constant time. Step 2, on a complete binary tree on n leaves would require $2 \lg n$ steps.
$T = 1 + 1 + 2 \lg n$. $P = 2n - 1 = O(n)$.

A segmented prefix (scan) computation consists of a sequence of disjoint prefix computations. Let the x_{ij} below take values from a set X and let $\oplus$ be an associative operator defined on the elements of set X. Then the segmented prefix computation for

$$x_{11}x_{12}\ldots x_{1k_1} \mid x_{21}x_{22}\ldots x_{2k_2} \mid \ldots \mid x_{m1}x_{m2}\ldots x_{mk_m} \mid$$

requires the computation of all

$$p_{ij} = x_{i1} \oplus x_{i2} \oplus \ldots \oplus x_{ij} \;\; \forall\, 1 \leq i \leq m,\; 1 \leq qj \leq k_i$$

In bried the segment separator | terminates one prefix operation and starts another one.
One way to deal with a segmented prefix computation in parallel is to extend $(X, \oplus)$ into $(X', \otimes)$ so that

$$X' = X \cup \{|\} \cup \{|\, x : x \in X\}$$

i.e. X' has more than twice the elements of X: it has all the elements of X, the segment separator | and a new element $|\, x$ which consists of the segment separator and x. The new operator $\otimes$ is associative if we define it as follows.

$\mid \otimes \mid = \mid$,	$\mid \otimes x = \mid x$,	$\mid \otimes \mid x = \mid x$,
$x \otimes \mid = \mid$,	$\mid x \otimes \mid = \mid$,	$x \otimes y = x \oplus y$
$\mid x \otimes y = \mid (x \oplus y)$	$x \otimes \mid y = \mid y$	$\mid x \otimes \mid y = \mid y$

Now, if the length of the segmented prefix formula is n we can assign n processors to solve the problem with parallel prefix in asymptotically the same time. Note that an element in X' requires for its representation no more than 2 extra bits of the storage size of an element of X. If an $\oplus$ computation takes $O(1)$ time so does an $\otimes$ computation.

(c) Copyright A. Gerbessiotis. All rights reserved.

Segmented Parallel Prefix: Example and Refinements

Example

$\{2,3\}\ \{1,7,2\}\ \{1,3,6\}$.

Create

$2\ \ 3\ \ |1\ \ 7\ \ 2|1\ \ 3\ \ 6$.

And result is:

$2\ \ 5|1\ \ 8\ \ 10|1\ \ 4\ \ 10$.

We can refine the previous algorithm as follows.

$\oplus$	b
a	$(a \oplus b)$

$\otimes$	b	$\|b$
a	$(a \oplus b)$	$\|b$
$\|a$	$\|(a \oplus b)$	$\|b$

(c) Copyright A. Gerbessiotis. All rights reserved.

Matrix Multiplication

A simple algorithm for multiplying two $n \times n$ matrices on a CREW PRAM with time complexity $T = O(\lg n)$ and $P = n^3$ follows. For convenience, processors are indexed as triples (i, j, k), where $i, j, k = 1, \ldots, n$. In the first step processor (i, j, k) concurrently reads a_{ij} and b_{jk} and performs the multiplication $a_{ij}b_{jk}$. In the following steps, for all i, k the results $(i, *, k)$ are combined, using the parallel sum algorithm to form $c_{ik} = \sum_j a_{ij}b_{jk}$. After $\lg n$ steps, the result c_{ik} is thus computed.

The same algorithm also works on the EREW PRAM with the same time and processor complexity. The first step of the CREW algorithm need to be changed only. We avoid concurrency by broadcasting element a_{ij} to processors $(i, j, *)$ using the broadcasting algorithm of the EREW PRAM in $O(\lg n)$ steps. Similarly, b_{jk} is broadcast to processors $(*, j, k)$.

The above algorithm also shows how an n-processor EREW PRAM can simulate an n-processor CREW PRAM with an $O(\lg n)$ slowdown.

```
                                          CREW    EREW
1. aij to all            (i,j,*) procs    O(1)    O(lgn)
   bjk to all            (*,j,k) procs    O(1)    O(lgn)
2. aij*bjk at            (i,j,k) proc     O(1)    O(1)
3. parallel sum aij *bjk (i,*,k) procs    O(lgn)  O(lgn)   n procs participate
             j

4. cik = sum aij*bjk                      O(1)    O(1)
           j


               3        3                 3
 T=O(lgn),P=O(n ) W=O( n  lgn)   W  = O(n )
                                  2
```

Problem. Let $X_1 \ldots, X_n$ be binary/boolean values. Find $X = X_1 \wedge X_2 \wedge \ldots \wedge X_n$.

The sequential problem accepts a $P = 1, T = O(n), W = O(n)$ direct solution.

An EREW PRAM algorithm solution for this problem works the same way as the PARALLEL SUM algorithm and its performance is $P = O(n), T = O(\lg n), W = O(n \lg n)$ along with the improvements in P and W mentioned for the PARALLEL SUM algorithm.

In the remainder we will investigate a CRCW PRAM algorithm. Let binary value X_i reside in the shared memory location i. We can find $X = X_1 \wedge X_2 \wedge \ldots \wedge X_n$ in constant time on a CRCW PRAM. Processor 1 first writes an 1 in shared memory cell 0. If $X_i = 0$, processor i writes a 0 in memory cell 0. The result X is then stored in this memory cell.

```
begin LOGICAL_AND (X_1 ... X_n)
1.   Proc 1   writes an    1   in cell 0.
2.   if  X_i = 0 processor i writes 0 into cell 0.
end LOGICAL_AND
```

The result stored in cell 0 is 1 (TRUE) unless a processor writes a 0 in cell 0; then one of the X_i is 0 (FALSE) and the result X should be FALSE, as it is.

Exercise 4 *Give an $O(1)$ CRCW algorithm for LOGICAL OR.*

Problem. Let $X_1 \ldots, X_N$ be n keys. Find $X = \max\{X_1, X_2, \ldots, X_N\}$.

The sequential problem accepts a $P = 1, T = O(N), W = O(N)$ direct solution.

An EREW PRAM algorithm solution for this problem works the same way as the PARALLEL SUM algorithm and its performance is $P = O(N), T = O(\lg N), W = O(N \lg N), W_2 = O(N)$ along with the improvements in P and W mentioned for the PARALLEL SUM algorithm.

In the remainder we will investigate a CRCW PRAM algorithm. Let binary value X_i reside in the local memory of processor i.

The CRCW PRAM algorithm MAX1 to be presented has performance $T = O(1)$, $P = O(N^2)$, and work $W_2 = W = O(N^2)$.

The second algorithm to be presented in the following pages utilizes what is called a doubly-logarithmic depth tree and achieves $T = O(\text{lglg}N)$, $P = O(N)$ and $W = W_2 = O(N \,\text{lglg}N)$.

The third algorithm is a combination of the EREW PRAM algorithm and the CRCW doubly-logarithmic depth tree-based algorithm and requires $T = O(\text{lglg}N)$, $P = O(N)$ and $W_2 = O(N)$.

(c) Copyright A. Gerbessiotis. All rights reserved.

begin MAX1 $(X_1 \ldots X_N)$
1. **in proc** (i,j) **if** $X_i \geq X_j$ **then** $x_{ij} = 1$;
2. **else** $x_{ij} = 0$;
3. $Y_i = x_{i1} \wedge \ldots \wedge x_{in}$;
4. Processor i reads Y_i ;
5. **if** $Y_i = 1$ processor i writes i into cell 0.

end MAX1

In the algorithm, we rename processors so that pair (i,j) could refer to processor $j \times n + i$. Variable Y_i is equal to 1 if and only if X_i is the maximum.

The CRCW PRAM algorithm MAX1 has performance $T = O(1)$, $P = O(N^2)$, and work $W_2 = W = O(N^2)$.

(c) Copyright A. Gerbessiotis. All rights reserved.

Doubly Logarithmic-Depth Tree

In preparation of algorithm MAX2 we introduce a **doubly logarithmic-depth tree.**
Let $N = 2^{2^n}$, for some integer n.
A *doubly logarithmic-depth* tree with N leaves is defined as follows.

(1) The root of the tree at level 0 has $2^{2^{n-1}} = N^{1/2}$ children in level 1.

(2) Each node at level 1 has $2^{2^{n-2}} = N^{1/2^2}$ children in level 2.

(3) Each node of level i has $2^{2^{n-(i+1)}} = N^{1/2^{i+1}}$ children in level $i+1$.

(4) Each node of level $n-1$ (the level before the last) has $2^{2^{n-n}} = N^{1/2^n} = 2$ children in level $n = \lg\lg N$.

(5) The nodes of level n are the **leaves** of the tree.

Some properties of a *doubly logarithmic-depth* tree are listed below.

(1) The **height** of the tree is $n = \lg\lg N$.

(2) A node of level i has $2^{2^{n-(i+1)}}$ children in level $i+1$.

(3) The TOTAL number of level i nodes is $2^{2^{n-1}} 2^{2^{n-2}} \ldots 2^{2^{n-i}} = 2^{2^n - 2^{n-i}}$.

(4) The Product $(2^{2^{n-i-1}})^2 \times 2^{2^n - 2^{n-i}}$ is $O(2^{2^n}) = O(N)$.

(c) Copyright A. Gerbessiotis. All rights reserved.

Algorithm MAX2

Algorithm MAX2 below achieves better work performance than MAX1: $T = O(\lg\lg N)$, $P = O(N)$, and $W = W_2 = O(N \lg\lg N)$.

Algorithm MAX2 works as follows.

(1) Initially, items X_i are on the N leaves of the tree.

(2) The root will hold the result at the completion of the algorithm.

(3) Processors are assigned to the nodes of the tree in some predetermined fashion.

(4) All nodes of the tree other than the leaves hold an UNDEFINED value in the beginning of the execution.

(5) If a node u at level i holds an UNDEFINED value and its M children hold some intermediate results M^2 processors are assigned to u to find the maximum of M numbers (the partial results of the children of u) using MAX1 in constant time. Node u then holds the computed maximum (and ceases to hold an UNDEFINED value).

```
begin MAX2  (X_1 ... X_N)
0.   The i-th cell contains X_i; N = 2^{2^n}.
1.   for  (i = n − 1; i ≥ 0; i − −) do
2        begin
3.           Assign to each node u of level i, (2^{2^{n−i−1}})^2 processors (i.e. the square of its children in level i + 1.) ;
4.           Use algorithm MAX1 and these processors to find the maximum
             of the values stored at the children of u and store the result at u;
5.       Node u ceases to hold an UNDEFINED.
6.       end
end MAX2
```

Algorithm MAX2 has $W = W_2 = O(N \lg\lg N)$. Algorithm MAX3 below has $W_2 = O(N)$. It uses MAX2 and the EREW PRAM algorithm as subroutines.

The EREW algorithm finds the SUM or the MAXIMUM of N numbers by working from the leaves to the root of a binary tree, ie $\lg N$ levels. If we stop the computation after i levels, we have $N/2^i$ partial results, each result being the MAXIMUM of 2^i numbers initially stored in the leavs of the subtree rooted at the partial result.

Max3 first runs the EREW PRAM algorithm for $i = \lg\lg\lg N$ levels so that a total of $N/\lg\lg N$ partial MAXIMA are computed.

Then it applies MAX2 where N in MAX2 is equal to the number of partial results ie $N/\lg\lg N$.

begin MAX3 $(X_1 \ldots X_N)$
0. The i-th cell contains X_i;
1 **begin**
2. Use the EREW PRAM algorithm on a complete binary tree on N leaves to reduce the original problem to computing the maximum of $N/\lg\lg N$ values (ie proceed from the leaves up to the nodes of level $\lg N - \lg\lg\lg N$);
3. Use MAX2 to find the maximum of the $N/\lg\lg N$ values of Step2;
4. **end**

end MAX3

Step 2 of algorithm MAX3 requires $T = O(\lg\lg\lg N)$, $P = O(N)$ and $O(N)$ work (number of comparisons is at most the number of edges of the tree).

Step 3 requires $P = O(N)$, $T = O(\lg\lg N)$ and total work $W = W_2 = O(N)$ by the analysis of MAX2.

Question. Is there a $p \leq n$ processor CRCW PRAM algorithm that finds the maximum of N keys faster than MAX2 or MAX3?

(c) Copyright A. Gerbessiotis. All rights reserved.

We introduce a special case of sorting n keys that are integers whose values are in the range of $[0..\lg n-1]$, where $\lg n$ we assume it is also an integer. This is the well-known problem of count-sort. Sorting n keys in the range $0..k-1$ requires sequential time $O(n+k)$. We show below that sorting n integer keys in the specified range with $P = n/\lg n$ processors can be done in time $T = O(\lg n)$ using $W = W_2 = \Theta(n)$. Note that count-sort does not sort in place i.e. we are going to use different arrays for input and output. Let $M[1..n-1]$ be the input and $N[1...n-1]$ the output arrays.

The idea is to assign $\lg n$ keys per processor; for example if keys are in $M[1..n-1]$, processor $i = 0, \ldots, p-1$ deals with keys $i \lg n + j + 1$, where $j = 0, \ldots, \lg n - 1$. The P processors collectively create an $P \times \lg n$ array C initialized to 0. We assign to each procesor a single row of the array. Thus initialization take $O(\lg n)$ steps to zero the entries of row i assigned to processor i. Entry (i, j) of the table would indicate how many keys with value j processor i is assigned to. This information can be collected easily: processor i scans its keys and for each key it updates the counters of row i of C. Total time is $O(\lg n)$. Then a parallel prefix operation is formed. It consists of the first column, the second column and so on the last column. The purpose is to count all the keys with values 0 before the ones with value 1, before those with value 2 and so on. Note that the prefix sequence is of length n. If we have n processors we can work it out in $O(\lg n)$ time. Now that we have only P processors we can invoke Brent's principle to do it in $O(\lg n)$ time as well but with $O(n)$ work. Note that during the prefix operation C is not overwritten; a new array D will hold the results. After the prefix operation if the entry that corresponded initially to the (i, j) element of C has value t, this means that processor i will store the keys assigned to it with value j to consecutive positions ending with memory location t of the output. Thus if for example we have that $C(i, j) = 3$, then N[t-2] N[t-1] and N[t] will hold the three keys with value j of processor i. Processor i, after D becomes available scans its keys and writes them into N as appropriately.

Thm. Sorting n keys in the range $[0..\lg n-1]$ with $P = n/\lg n$ processors can be done in time $T = O(\lg n)$ and work $W = W_2 = \Theta(n)$.

Applying this theorem t times (i.e. use t rounds of count-sort to obtain a radix-sort algorithm) the following is derived.

Thm. Sorting n keys in the range $[0..\lg^t n-1]$ with $P = n/\lg n$ processors can be done in time $T = O(t \lg n)$ and work $W = W_2 = \Theta(tn)$.

(c) Copyright A. Gerbessiotis. All rights reserved.

Parallel Count-Sort: Sequential algorithm

We show the corresponding sequential algorithm which also indicates the "sequential work" that each one of the P processors will perform on its own $\lg n$ assigned keys.

```
 Count-Sort(M[1..n],N[1..n],n,k)
   // Initialize Counter array C
1. for(i=0;i<k;i++)   C[i]=0;
3. for(j=1;j<=n;j++)  C[M[j]] ++;
4. D[0]=C[0]; //Note that the extra D is not required; we can reuse C
5. for(i=1;i<k;i++)
6.   D[i] = C[i]+D[i-1];
7. for(j=n;j>=1;j--) {
8.    N[D[M[j]]]=M[j];
9.    D[M[j]]--;
   }
```

(c) Copyright A. Gerbessiotis. CS 667-101 : Spring 2012 . All rights reserved.

Parallel Count-Sort: An example

```
Step 1:  Split keys and initialize C; Count keys and update C
M: 0 2 2 0   1 3 0 0   1 3 0 3  2 1 3 2          (n=16,lgn=4)
P:    0         1         2        3


C:
  p   0  1  2 3
   0  2  0  2 0 <<< Processor zero counted two 0's and two 2's for its elements
   1  2  1  0 1
   2  1  1  0 2
   3  0  1  2 1

Step 3:  For a prefix: take first column  second colum  ...  lgn-th column

Prefix: Input:     2 2 1 0   0 1 1 1   2  0  0 2   0  1 2  1    : C array info
       Output:     2 4 5 5   5 6 7 8  10 10 10 12 12 13 15 16   : D array info

The length of the prefix is n.  n processors can do it in O(lg{n}) time.
n/lgn processors in O(lg{n}) time and O(n) work.
                   2 4 5 5   5 6 7 8  10 10 10 12 12 13 15 16   : D array info
                   *         *         *          *
Processor 0 recovers four numbers 2,5,10,12 the last position occupied by
                              a   0,1, 2, 3 if it has them
It has 2 0's that will occupy positions 1 and 2  of the output
and    2 2's that will occupy positions 9 and 10 of the output
He copies them into N as dictated by array D (observe sequential algorithm)
```

(c) Copyright A. Gerbessiotis. All rights reserved.

A matching Lower bound and Turan's Theorem

Definition. On an undirected graph $G = (V, E)$, an **independent set** is a set of vertices such that no two vertices are connected by an edge.

Definition. On an undirected graph $G = (V, E)$, a **clique** is a set of vertices such any two vertices are connected by an edge.

A clique on n vertices is the "complement" of an independent set on the same n vertices. The complete graph on n vertices is a clique on n vertices by default. A clique of n vertices has $n(n-1)/2$ edges.

Theorem (Turan) Let $G = (V, E)$ be an undirected graph, where $|V| = n$ and $|E| = m$. Graph G has an independent set of size at least $n^2/(2m + n)$.

Another formulation of Turan's Theorem is the following one.

Theorem (Turan): Second version Let $G = (V, E)$ be an undirected graph, where $|V| = n$ and $|E| = m$. If graph G has no p clique then it has at most $(1 - 1/(p-1))n^2/2$ edges. **Proof.** If $n \leq p - 1$ then G does not have obviously a p-clique and

G has at most $n(n-1)/2$ edges. It is obvious that $n(n-1)/2 \leq (1 - 1/(p-1))n^2/2$ by elementary calculation.

Thus the interesting case left is $n \geq p$. If graph G has the maximum number of edges but does not have a p-clique it must have a $(p-1)$-clique. This is because otherwise we could add edges to G to create such a $(p-1)$-clique; this would contradict the maximality of edges of G. Call C a $(p-1)$-clique of G. C has $(p-1)(p-2)/2$ edges. Call G' the graph G without C, i.e. $G' = G - C$. Graph G' has m' edges. Let k be the number of edges that go from G' to C. By induction on G' we have that $m' \leq (1 - 1/(p-1))(n - p + 1)^2/2$. Since G does not have a p-clique every vertex of G' is connected to at most $p - 2$ vertices of C (since if it were connected to all the vertices of C a p-clique would have been formed). Thus $k \leq (p-2)(n-p+1)$.

Therefore the number of edges m of G is

$$m \leq (p-1)(p-2)/2 + (1 - 1/(p-1))(n-p+1)^2/2 + (p-2)(n-p+1) \leq (1 - 1/(p-1))n^2/2.$$

(c) Copyright A. Gerbessiotis. All rights reserved.

If we solve this inequality for p we get

$$p - 1 \geq \frac{n^2}{n^2 - 2m}$$

Therefore an equivalent formulation is that G has a $p-1$ clique of size at least $\frac{n^2}{n^2-2m}$.

Now take the complementary graph (where an edge becomes a non-edge and a non-edge becomes an edge). The complement of G has $N = n$ vertices and $M = n(n-1)/2 - m$ edges. From the latter we get that $2m = n(n-1) - 2M = n^2 - n - 2M$. A $p-1$ clique in G becomes an independent set in its complement whose size is at least $\frac{n^2}{n^2-2m} = \frac{N^2}{N^2-2m} = \frac{N^2}{N^2-n^2+n+2M} = \frac{N^2}{N+2M}$, noting that $N = n$. This latter bound is the expression in the first version of Turan's theorem.

An easy corollary is that a graph with n/k vertices, $k \geq 1$, and n edges has an independent set of size at least $n/4k^2$.

We are going to show some lower bounds on finding the MAXIMUM of n keys. The model of computation we are going to use is the decision tree model, and in fact the **parallel decision tree model** where we allow p processors to work at any time step each one performing a single comparison between two keys. The decision tree model in the case $p = 1$ was used to prove the lower bound $\Omega(n \lg n)$ for comparison-based sorting. Note that this model deals with **information gathering** to compute the MAXIMUM. One also needs to process this information to derive the MAXIMUM. The model assume that processing is **for free**. We can do so because we plan to establish a lower-bound (minimum possible running time) not to establish an algorithm for realizing this bound.

We know from the sequential setting that finding the MAXIMUM of n keys requires at least $n - 1$ comparisons. **Lemma 1.** MAX can be found in ONE parallel step with $n(n-1)/2$ processors.

Proof. n keys allow $n(n-1)/2$ pairs and thus comparisons to be realized to obtain all the information required to find the MAXIMUM (one needs to do some additional processing eg. establishing the rank but that it is for free!). If we have that many processors each one responsible for one comparison, this concludes the proof.

Lemma 2. In order to find MAX in ONE parallel step we need at least $n(n-1)/2$ processors.

Proof. In order to prove a lower bound, we use a proof by contradiction. Suppose we can do it with one fewer processor $P = n(n-1)/2 - 1$. Since there are $n(n-1)/2$ pairs of keys to compares, one such pair is not compared. Call the keys of the pair x, y. What we are going to show, by playing the role of an adversary, is that we can set up the values of the n input keys so that the missing comparison of the x, y is the crucial one to establish the maximum. We thus play the role of an adversary whose only mission is to make the algorithm that uses P processors to fail. To do so, we set the results of the comparisons in such a way that x, y are the MAX and SECOND MAX keys. Thus the comparison between x, y (that is not being performed) is CRUCIAL in determining the MAXIMUM of the n keys. Since it is not performed we cannot find the MAX with P processors. Contradiction is established.

Question. What can you prove about SECOND MAX?

(c) Copyright A. Gerbessiotis. All rights reserved.

Some preliminry results

An interesting question is how many processors one needs to use to find the MAX of n keys not in one parallel step but in two parallel steps.

Problem 1. Find MAX in two steps with $O(n^{3/2})$ processors.

Hint. Split n keys into groups of $\sqrt{n}$. Compute MAX of each group in first step, and MAX of MAXes in second step.

Problem 2. Find MAX in two steps with $O(n^{4/3})$ processors.

Hint. Optimize the splitting. $\sqrt{n}$ might not be optimal.

Problem 3. Show that MAX in two steps requires $\Omega(n^{4/3})$ processors.

The Proof is a repetition of the arguments of the proof of Theorem 1.

(c) Copyright A. Gerbessiotis. CS 667-101 : Spring 2012 . All rights reserved.

Theorem 1 (Valiant) Computing the maximum of n keys requires at least lglg$n$ parallel steps with $p \leq n$ processors.
Proof. (by induction) It is proved by what we call **an adversary argument** through induction. An **adversary** for this problem is allowed to choose the input keys by modifying their values in such a way so as to force the algorithm to run for at least lglg$n$ steps. These modifications should not invalidate, however, the operations of the algorithm already performed.

View step i as a graph $G_i = (V_i, E_i)$. The vertices are the keys and the edges are the comparisons performed at step i.

Consider initially graph G which becomes G_1. Since any max finding algorithm can perform no mores than $p \leq n$ comparisons at any time step, $|E_1| \leq n$ and $|V_1| = n$. So by Turan's theorem G_1 has an independent set of size $n/4$. Call this set of keys I_1. Consider now the computation where all the I_1 keys are the larger than anything in $V_1 - I_1$. In this case every node in I_1 wins in the comparisons performed and is a candidate for the maximum. Since I_1 is an independent set we have no information on the relative order of the keys in I_1.

Consider now G_2 and take the intersection of G_2 and I_1. It has $\leq n$ edges (since $p \leq n$ can be performed at a time) and at least $n/4$ vertices (lower bound for I_1 size. So the graph G_2 intersected by I_1 has an independent set of size at least $n/64$ and call it I_2. Repeating the same thing for G_i and I_{i-1} we end up with an independent set of size $n/2^{2^{i+1}-2}$. So if we repeat this procedure about $\Omega(\lg \lg n)$ times the independent set will drop below 2 and the maximum will be established. This takes however $\Omega(\lg \lg n)$ parallel steps.

(c) Copyright A. Gerbessiotis. All rights reserved.

Let us prove Problem 3.

Proof (Problem 3). We set up a graph, just as in Theorem 1, with n vertices. Let us have p processors. In one step they can force p comparisons. By Turan's theorem the graph on n vertices and p edges must have an independent set of size at least $k = n^2/(n + 2p)$. We can set the values of the keys or equivalently determine the output of the comparisons so that the keys of the named independent set are all candidates for the MAX. Since they form an independent set none has been compared to any other key of the set. Thus in the second round we can find the MAX among these k keys in $k(k-1)/2$ comparisons/processors using Lemma 1, if and only if we can afford to do so i.e. $p \geq k(k-1)/2$, which leads to $p = \Omega(n^{4/3})$. □

Problem 4. Is there an algorithm that finds the MAX in $O(\lg \lg n)$ parallel steps using n processors? What is the work of the algorithm?

Long Hint. Consider n keys. Splits into $n/3$ groups of 3 keys each. For each group we can detemine the MAX in $3(3-1)/2 = 3$ comparisons using 3 procs per group. Total number of processors used is $n/3 \cdot 3 = n$. Thus we are left with determining the max of $n/3$ keys.

Take the $n/3$ Maxima, and split them into groups of 7. We have $n/(3 \cdot 7)$ groups of 7 keys. Each group requires $7(7-1)/2 = 21$ processors to find the MAX of the group. Total processors used is $n/21 \cdot 21 = n$, that we can afford to. Thus after the second second it suffices to find the max of $n/21$ keys to determine the MAX of the original n keys.

How do we split the $n/21$ keys next? What is the pattern?

Say we at some point we end up having n/s keys. We split them into groups of t so that $t(t-1)/2$ comparisons/processors per group. You can fill in the details to show that this way $\lg \lg n$ can be achieved.

(c) Copyright A. Gerbessiotis. All rights reserved.

Theorem PSum. Let the minimum number of operations for computing an expression E be $m \geq p$. Then evaluation of E in parallel using p processors requires time at least

$$T \geq \left\lceil \frac{(m+1-2^{\lceil \lg p \rceil})}{p} \right\rceil + \lceil \lg p \rceil$$

if $m \geq 2^{\lceil \lg p \rceil}$, and

$$T \geq \lceil \lg (m+1) \rceil$$

otherwise.

Proof. We assume that each processor computes at every step a simple expression $z = x\ o\ y$, where $o \in \{+, -, /, *\}$.

In the last step of the evaluation of E, the T-th step, a single operation like $E = z_1 o z_2$, is performed on a single processor. Then, in step $T-1$, operations $z_1 = y_1 o y_2$ and $z_2 = y_3 o y_4$ may be performed, i.e. at most two operations are performed in parallel; at most two processors can be employed to compute the arguments z_1, z_2 of the operation of the T-th step. Then in the $T-i$-th step at most $\min(p, 2^i)$ processors can actively take part in computing E.

This means that in the $\lceil \lg p \rceil$ steps ending with T i.e. the T-th, $(T-1)$-st all the way down to $(T-(\lceil \lg p \rceil - 1))$-th step, the total number of operations performed is $\sum_i^{\lceil \lg p \rceil - 1} 2^i = 2^{\lceil \lg p \rceil} - 1$.

We distinguish two cases.

Case 1. If $m > 2^{\lceil \lg p \rceil} - 1$, then $T > \lceil \lg p \rceil$. In the remaining steps, which are $T - \lceil \lg p \rceil$, p processors can perform p operations simultaneously. Thus a total of $p(T - \lceil \lg p \rceil) + 2^{\lceil \lg p \rceil} - 1$ operations can be performed in the T step for this case. Therefore $m \leq p(T - \lceil \lg p \rceil) + 2^{\lceil \lg p \rceil} - 1$. and the result follows, since T (number of parallel steps) is an integer.

Case 2. If $T \leq \lceil \lg p \rceil$ or $m \leq 2^{\lceil \lg p \rceil} - 1$, then it follows that complete parallelism can not be achieved (i.e. some processors will remain idle). Then, $m \leq 2^T - 1$ and consequently, $T \geq \lceil \lg (m+1) \rceil$. This completes the other case. ■.

Theorem M-P. Let $n = 2m+1$ be odd. Let $f(x) = a_n x^n + \ldots + a_1 x + a_0$, and let $p \lg n \leq n$. Then $f(x)$ can be evaluated in $T = 2n/p + \lg p + c$ multiplications and additions, where c is a constant less than 4. The constant c absorbs the contributions of $\lceil \lg p \rceil$ factors that are simplified as $\lg p$.

Proof. Phase 1. We keep processors busy by performing a Horner rule computation and a parallel prefix computation $f(x) = ((a_n x + a_{n-1})x + \ldots$ by working as follows.

(a) We first compute in parallel $b_0 = a_0 + a_1 x,\ b_1 = a_2 + a_3 x,\ \ldots, b_m = a_{n-1} + a_n x$. The number of terms we compute in parallel is $(n+1)/2 = m+1$, and we thus need to perform a total of $m+1$ A's and $m+1$'s M's to complete this subtask of Phase 1. The total parallel running time is $2(m+1)/p = (n+1)/p$.

(b) We then compute $t = x^2$ and then by performing a parallel prefix operation we compute t, t^2, t^3, t^p. This requires a total of p M's. Since these can be performed in parallel no processor performs more than $\lceil \lg p \rceil = \lg p + 1$ M's (and the 1 is absorbed in c hereafter).

The two steps can be interleaved. Between the two steps, the computation of step (a) takes more time. Therefore the total time for Phase 1 is $(n+1)/p$.

Phase 2. Let $b_0 = a_0 + a_1 x$, $b_1 = a_2 + a_3 x \ldots, b_m = a_{n-1} + a_n x$.

Then, for $t = x^2$, $f(x)$ can be rewritten as follows.

$$f(x) = b_0 + b_1 x^2 + b_2 x^4 + \ldots + b_m x^{2m} = b_0 + b_1 t + b_2 t^2 + \ldots + b_m t^m.$$

We note that the m-th term $b_m t^m = (a_{n-1} + a_n x)t^m = (a_{n-1} + a_n x)x^{2m} = (a_{n-1}x^{n-1} + a_n x^n)$ contributes the two highest degree terms of $f(x)$. Then, we can rewrite $f(x)$ as follows.

$$\begin{aligned} f(x) &= t^0(b_0 + b_p t^p + b_{2p} t^{2p} + \ldots + b_{k_0 p} t^{k_0 p}) + \\ &\quad t^1(b_1 + b_{p+1} t^p + b_{2p+1} t^{2p} + \ldots + b_{k_1 p+1} t^{k_1 p}) + \\ &\ldots \\ &\quad t^{p-1}(b_{p-1} + b_{2p-1} t^p + b_{3p-1} t^{2p} + \ldots + b_{k_{p-1} p+1} t^{k_{p-1} p}) \\ &= t^0 g_0(t^p) + t^1 g_1(t^p) + \ldots + t^i g_i(t^p) + \ldots + t^{p-1} g_{p-1}(t^p) \end{aligned}$$

where $g_i(t^p)t^i$ is the sum of all $b_j t^j$ such that $j \equiv i \bmod p$. Each $g_i(t^p)$ is a polynomial of degree $k_i \leq m/p$ in t^p.

Phase 2 then consists of the following three steps (a), (b) and (c).

Munro-Paterson method (continued)

(a) Assign processor i to computing $g_i(t^p)$ for all $0 \leq i \leq p-1$. t^p was computed in step (b) of Phase 1. Processor i then evaluates polynomial $g_i(t^p)$ given by

$$g_i(t^p) = b_i + b_{p+i}(t^p)^1 + b_{2p+i}(t^p)^2 + \ldots + b_{k_ip+i}(t^p)^{k_i}$$

The polynomial is of degree an most m/p in t^p, and thus its evaluation requires m/p A's and m/p M's by using Horner's rule for a total of $2m/p = (n-1)/p$.

(b) Compute on processor i, $g_i(t^p)t^i$, in parallel, by performing one multiplication per processor noting that all the t^i's have become available since the parallel prefix of Phase 1.

(c) The p results $g_i(t^p)t^i$ of step (b) are summed up in $\lg p$ steps to evaluate f at x. Using a parallel sum no processor performs more than $\lceil \lg p \rceil$ A's for a total of $\lg p$ parallel A's (and one more unit contribution to c).

Total running time of all steps is $(n+1)/p + (n-1)/p + \lg p + 1 + d = 2n/p + \lg p + c$. Note that d is no more than 2, and thus c is no more than 3.

(c) Copyright A. Gerbessiotis. All rights reserved.