

# Parallel Fast Fourier Transform


(PSC §3.4)



Parallel FFT

# Data distributions for butterflies of FFT

*butterfly distance*



- ▶  $n, p$  must be powers of two with  $p < n$ .
- ▶ In stage  $k$ , component pair  $(x_j, x_{j+k/2})$  at distance  $k/2$  is combined.
- ▶ Block distribution works for  $k = 2, 4, \dots, n/p$ .
- ▶ Cyclic distribution works for  $k = 2p, 4p, \dots, n$ .



## Block distribution works for small butterflies

Let  $n = 8, p = 2$ . In stage  $k = 2$ , the vector  $\mathbf{x}$  is multiplied by

$$I_4 \otimes B_2 = \begin{bmatrix} 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & -1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & -1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & -1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & -1 \end{bmatrix}.$$

- ▶ The first two butterfly blocks  $x(0: 1)$ ,  $x(2: 3)$  are contained in processor block  $x(0: 3)$ .
- ▶ The last two butterfly blocks  $x(4: 5)$ ,  $x(6: 7)$  are contained in processor block  $x(4: 7)$ .

## Cyclic distribution works for large butterflies

In stage  $k = 8$ , the vector  $\mathbf{x}$  is multiplied by

$$I_1 \otimes B_8 = B_8 = \begin{bmatrix} 1 & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & \cdot & \omega & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \omega^2 & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \omega^3 \\ 1 & \cdot & \cdot & \cdot & -1 & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & \cdot & -\omega & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot & \cdot & \cdot & -\omega^2 & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & -\omega^3 \end{bmatrix},$$

where  $\omega = \omega_8 = e^{-\pi i/4} = \frac{1}{2}\sqrt{2} - \frac{1}{2}\sqrt{2}i$ .

- ▶ The pairs  $(x_0, x_4)$  and  $(x_2, x_6)$  are combined on  $P(0)$ .
- ▶ The pairs  $(x_1, x_5)$  and  $(x_3, x_7)$  are combined on  $P(1)$ .



## Parallelisation strategy: use different distributions

- ▶ At the start, for  $k \leq n/p$ , we use the block distribution.
- ▶ Near the end, for  $k \geq 2p$ , the cyclic distribution.
- ▶ This suffices if  $p \leq n/p$ , i.e.  $p \leq \sqrt{n}$ . For example:  $p \leq 32$  for  $n = 1024$ .
- ▶ If  $p > \sqrt{n}$ , we need an **intermediate distribution**, a generalisation of the block and cyclic distribution.
- ▶ Split the vector into **blocks**. Each block is owned by a **group of processors** and is distributed by the **cyclic** distribution over the processors of that group.



# Group-cyclic distribution

- ▶ Let  $c$  be fixed such that  $1 \leq c \leq p$  and  $p \bmod c = 0$ . The **group-cyclic distribution with cycle  $c$**  is defined by


$$x_j \mapsto P((j \operatorname{div} \left\lceil \frac{cn}{p} \right\rceil)c + (j \bmod \left\lceil \frac{cn}{p} \right\rceil) \bmod c).$$

- ▶  $c$  is the number of processors in a group and  $\left\lceil \frac{cn}{p} \right\rceil = \left\lceil \frac{n}{p/c} \right\rceil$  is the size of a block owned by a group.
- ▶ If  $n \bmod p = 0$ , as happens in the FFT, this reduces to

$$x_j \mapsto P((j \operatorname{div} \frac{cn}{p})c + j \bmod c).$$

- ▶ For  $c = 1$ , we get the **block distribution**.  
For  $c = p$ , we get the **cyclic distribution**.

## From block to cyclic distribution



$n = 8$ ,  $p = 4$ , so that  $p > \sqrt{n}$ .

In (b), we have  $p/c = 2$  groups of  $c = 2$  processors.



## Global and local indices

- ▶  $n, p$  and hence  $c$  are powers of two. We have  $1 \leq c < \frac{cn}{p}$ .
- ▶ Thus, we can write the **global index**  $j$  as

$$j = j_2 \frac{cn}{p} + j_1 c + j_0,$$

where  $0 \leq j_0 < c$  and  $0 \leq j_1 < n/p$ .

- ▶ The processor that owns component  $x_j$  is

$$P((j \operatorname{div} \frac{cn}{p})c + j \operatorname{mod} c) = P(j_2 c + j_0).$$

- ▶ Processors in the same group have the same  $j_2$ , but  $j_0$  differs.
- ▶ We obtain the **local index**  $j$  by ordering the local components by increasing global index  $j$ , so that  $j = j_1$ .





# Which operations are local?

Butterfly operation on  $(x_j, x_{j+k/2})$  is **local** if

- ▶  $x_j, x_{j+k/2}$  are in the **same group**, i.e.  $k \leq \frac{cn}{p}$ ;
- ▶ distance  $k/2$  is a **multiple of  $c$** , i.e.  $k \geq 2c$ .

We can use the group-cyclic distribution with cycle  $c$  for


$$2c \leq k \leq \frac{n}{p}c.$$

**Outline of algorithm:**

- ▶ start with  $c = 1$ , perform stages  $k = 2, 4, \dots, n/p$ ;
- ▶ increase  $c$  to  $c = n/p$ , perform stages  $k = 2n/p, 4n/p, \dots, (n/p)^2$ ;
- ▶ ...
- ▶ finish with  $c = p$ , instead of  $c = (n/p)^t \geq p$ .



Warning: difficult slides ahead



# Parallel unordered FFT

```
{ distr(x) = block } k := 2; c := 1;
while k ≤ n do
(0)   j0 := s mod c; j2 := s div c;
      while k ≤  $\frac{n}{p}c$  do
        nblocks :=  $\frac{nc}{kp}$ ; { n/k butterfly blocks, p/c groups }
        for r := j2 · nblocks to (j2 + 1) · nblocks - 1 do
          { Compute part of x(rk: (r + 1)k - 1) }
          for j := j0 to  $\frac{k}{2} - 1$  step c do
            τ := ωkj xrk+j+k/2;
            xrk+j+k/2 := xrk+j - τ;
            xrk+j := xrk+j + τ;
          k := 2k;
```



# Parallel unordered FFT


```
{ distr(x) = block } k := 2; c := 1;
while k ≤ n do
(0)   j0 := s mod c; j2 := s div c;
       while k ≤  $\frac{n}{p}c$  do
           nblocks :=  $\frac{nc}{kp}$ ; { n/k butterfly blocks, p/c groups }
           for r := j2 · nblocks to (j2 + 1) · nblocks - 1 do
               { Compute part of x(rk : (r + 1)k - 1) }
               for j := j0 to  $\frac{k}{2} - 1$  step c do
                   τ := ωkj xrk+j+k/2;
                   xrk+j+k/2 := xrk+j - τ;
                   xrk+j := xrk+j + τ;

           k := 2k;
       if c < p then
           c0 := c; c := min( $\frac{n}{p}c, p$ );
(1)   redistr(x, n, p, c0, c);
       { distr(x) = cyclic }
```



# Parallel bit reversal

$$p = 8, n = 256$$



Start in cyclic distribution for index  $j$  with **local bit reversal**.  
Then **swap the data** between processors  $P(s)$  and  $P(\rho_p(s))$ .  
We end up in the block distribution for the reversed index  $\rho_n(j)$ .

# Postponing the processor swaps

- ▶ The distribution just before the swaps is the **block distribution with bit-reversed processor numbering**.
- ▶ All processors perform the same operations in FFT stages  $k = 2, 4, \dots, n/p$ , multiplying local blocks of  $\mathbf{x}$  by  $B_k$ .
- ▶ **I'll scratch your back if you scratch mine**: processors perform the work of their partner.
- ▶ The data swap can be postponed until the first redistribution, immediately after stage  $k = n/p$ .
- ▶ **Buy 2, Pay 1**: two permutations can be done at the cost of one by combining them. Hence **no extra communication** is incurred by the data swaps.



## Redistribution with possible proc-number reversal

*input:*  $\mathbf{x}$  : vector of length  $n = 2^m$ ,  
distr( $\mathbf{x}$ ) = group-cyclic with cycle  $c_0$  over  $p = 2^q$  procs.  
If *rev* is true, processor numbering is bit-reversed.

*output:* distr( $\mathbf{x}$ ) = group-cyclic with cycle  $c_1$ .

*call:* redistr( $\mathbf{x}, n, p, c_0, c_1, rev$ );

(1) **if** *rev* **then**  
     $j_0 := \rho_p(s) \bmod c_0$ ;  
     $j_2 := \rho_p(s) \operatorname{div} c_0$ ;  
**else**  
     $j_0 := s \bmod c_0$ ;  
     $j_2 := s \operatorname{div} c_0$ ;

**for**  $j := j_2 \frac{c_0 n}{p} + j_0$  **to**  $(j_2 + 1) \frac{c_0 n}{p} - 1$  **step**  $c_0$  **do**  
     $dest := (j \operatorname{div} \frac{c_1 n}{p}) c_1 + j \bmod c_1$ ;  
    put  $x_j$  in  $P(dest)$ ;



## Last iteration of main loop

- ▶ The last iteration is determined by the **smallest integer**  $t$  such that  $(n/p)^t \geq p$ .
- ▶ The cycles of the iterations are  $c = (n/p)^0, (n/p)^1, \dots, (n/p)^{t-1}, p$ .
- ▶ The **total number of iterations** is therefore  $t + 1$ .
- ▶ Since

$$\begin{aligned}(n/p)^t \geq p &\iff n^t \geq p^{t+1} \iff 2^{mt} \geq 2^{q(t+1)} \\ &\iff mt \geq q(t+1) \iff mt - qt \geq q \\ &\iff t \geq \frac{q}{m-q},\end{aligned}$$

it follows that

$$t = \left\lceil \frac{q}{m-q} \right\rceil.$$



## BSP cost

- ▶ Every iteration has a computation superstep and a communication superstep, except the last, which has no data redistribution. Therefore,

$$T_{\text{sync}} = (2t + 1)l = \left( 2 \left\lceil \frac{q}{m - q} \right\rceil + 1 \right) l.$$

- ▶ Every redistribution moves at most **all the local data** in and out, i.e.,  $n/p$  complex numbers, or  $2n/p$  real data words. Therefore,

$$T_{\text{comm}} = t \cdot \frac{2n}{p} g = \left\lceil \frac{q}{m - q} \right\rceil \cdot \frac{2n}{p} g.$$

- ▶ Look mama, without counting!

$$T_{\text{comp}} = (5n \log_2 n) / p.$$



# Summary

- ▶ We have used **different distributions** in different parts of the algorithm, trying to make our operations local.
- ▶ The algorithm **starts and finishes** in the cyclic distribution.
- ▶ If we split a vector into  $p/c$  **blocks** and distribute each block over  $c$  processors by the **cyclic** distribution, then we obtain the **group-cyclic distribution** with cycle  $c$ .
- ▶ The total BSP cost of the parallel FFT algorithm is

$$T_{\text{FFT}} = \frac{5n \log_2 n}{p} + 2 \cdot \left\lceil \frac{\log_2 p}{\log_2(n/p)} \right\rceil \cdot \frac{n}{p} g + \left( 2 \left\lceil \frac{\log_2 p}{\log_2(n/p)} \right\rceil + 1 \right) l.$$

- ▶ For practical  $p$ , we need only one data redistribution:

$$T_{\text{FFT}, 1 < p \leq \sqrt{n}} = \frac{5n \log_2 n}{p} + 2 \frac{n}{p} g + 3l.$$

