

ROUTING ?
Sorting?
Image Processing?
Sparse Matrices?

Reconfigurable
Meshes !



P Å R C



Reconfigurable architectures


- FPGAs
- reconfigurable multibus
- reconfigurable networks (Transputers, PVM)
- **dynamically reconfigurable mesh**

Aim:

efficiency

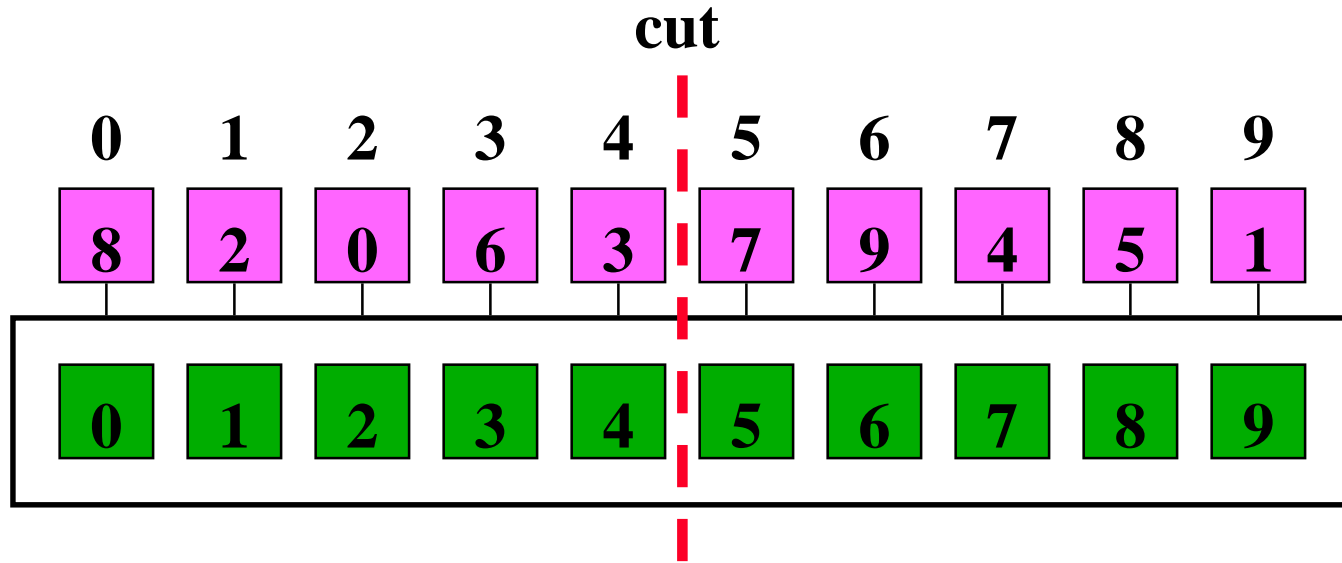
special purpose --> general purpose architectures

- 1.) Motivation for the reconfigurable mesh
- 2.) Routing (and sorting):
 - better than PRAM
 - better than mesh
- 3.) Image processing
- 4.) Sparse matrix multiplication
- 5.) Bounded bus length



P A R C

PRAM



diameter $O(1)$


bisection width $\Omega(n)$

EREW

CRCW

PARC

Mesh/Torus




2D mesh

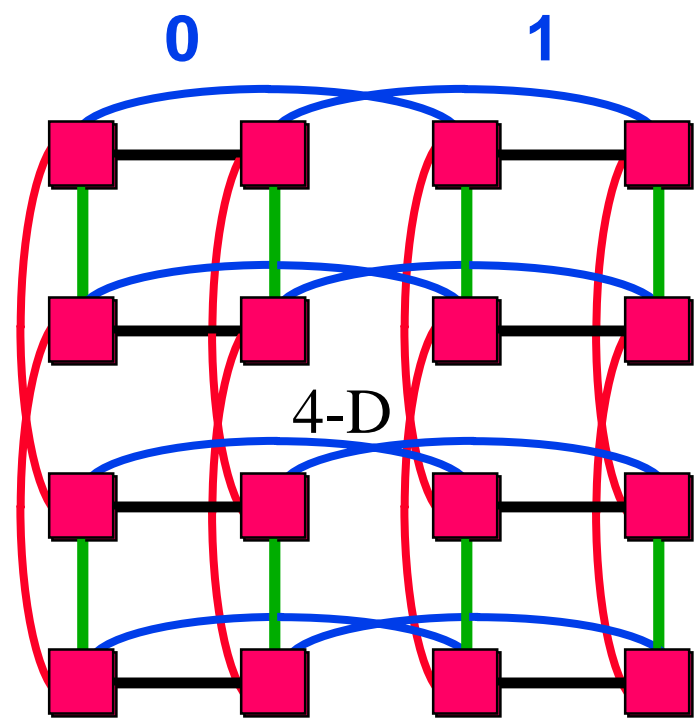
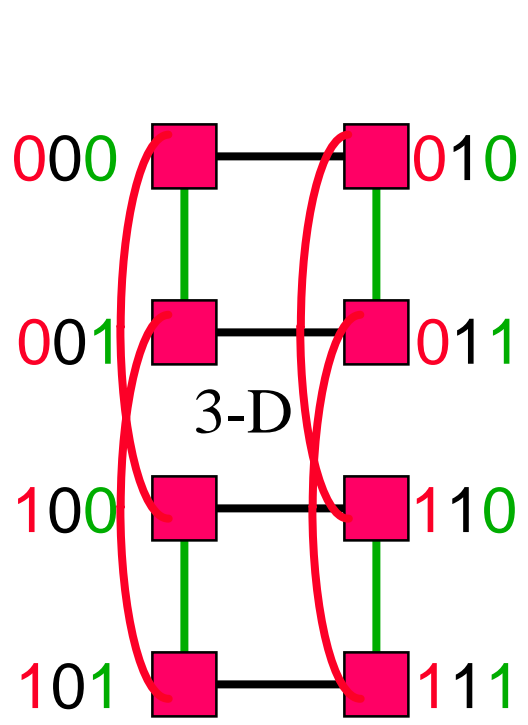
Diameter $\Theta(\sqrt{n})$
bisection width $\Theta(\sqrt{n})$

P A R C

Hypercube




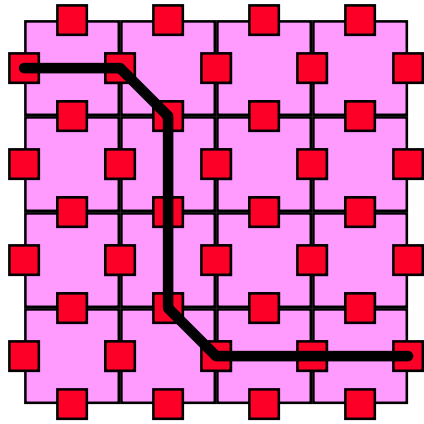
diameter $O(\log n)$
bisection width $\Theta(n)$



P A R C

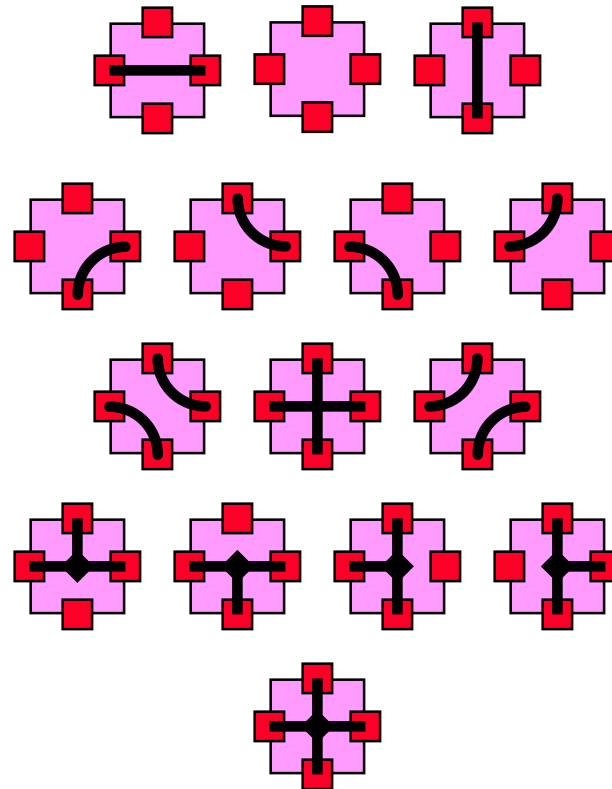
reconfigurable mesh

reconfigurable mesh = mesh + interior connections



low cost


diameter 1 !!



15 positions

P A R C

global OR




Time: $O(1)$ on RM

-- $\Omega(\log n)$ on EREW-PRAM

P A R C

Prefix sum

0 1 1 0 1 0 0 1 1 1



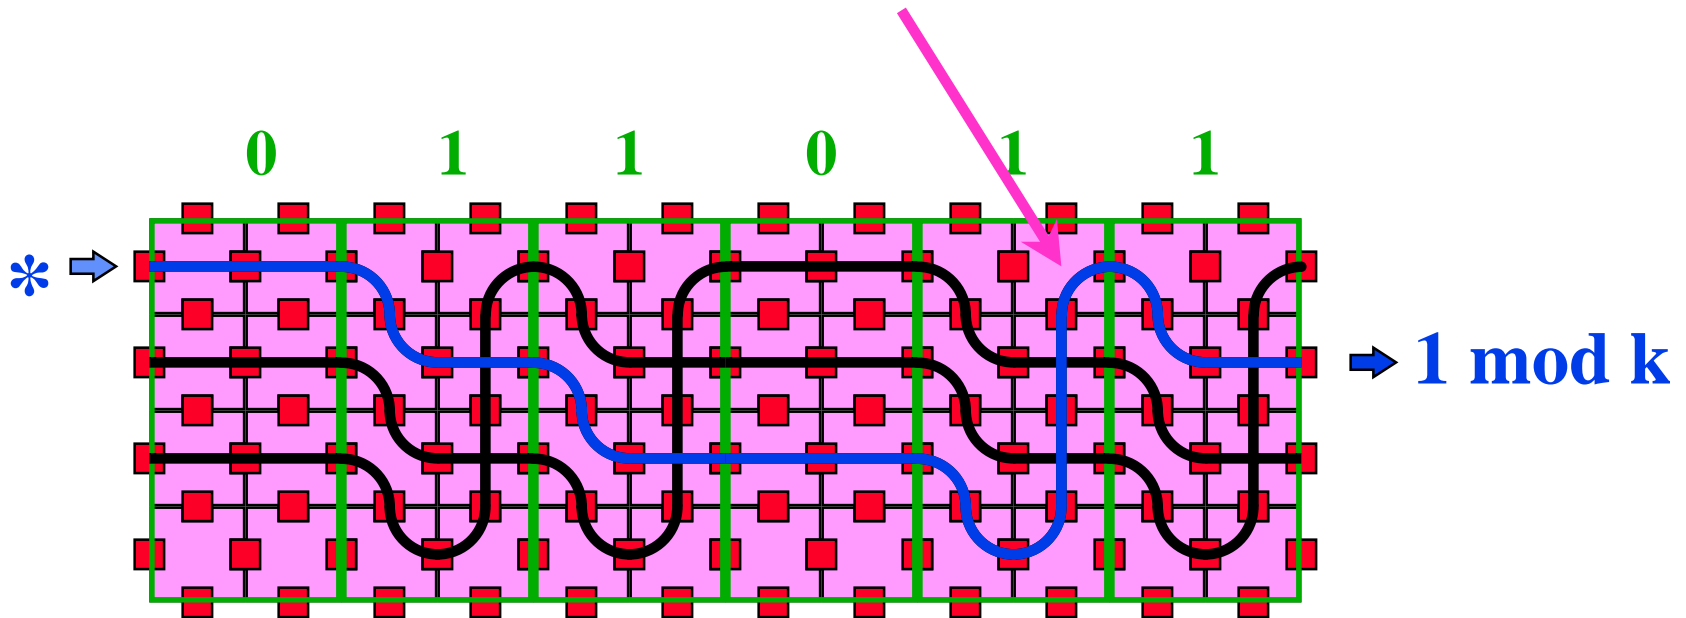
Time : $O(1)$
Area: $\Omega(n \times n)$

Fast but
expensive

P A R C

modulo k^2 counter (ranking)


- 2 digit numbers to the basis of k represent all numbers smaller than k^2 .
- 1.) determine $x \bmod k$ (=lsd)
- 2.) count number of “wraps” (=msd).



--> modulo k^2 counting in 2 steps on a $k \times k^2$ array

P A R C

enumeration / prefix sum




time: $O(\log n)$

wire efficiency ! -- (compared with tree)
1/2 number of processors

P A R C

permutation routing - 2 steps





$n \times n$

2 steps !!!



P A R C

Kunde's all-to-all mapping



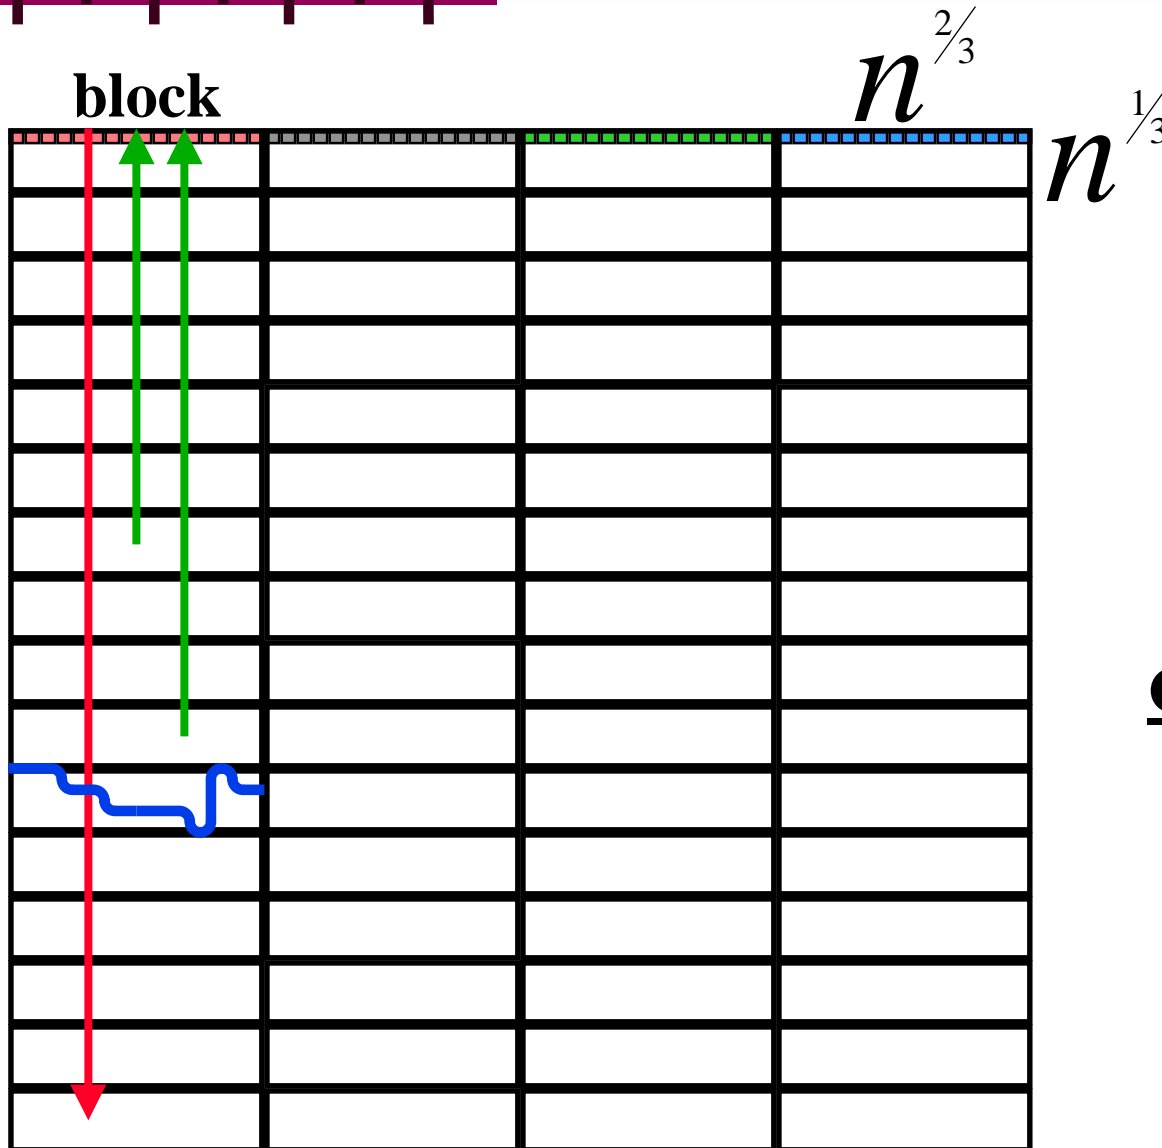
Sorting:

- sort blocks
- all-to-all (columns)
- sort blocks
- all-to-all (rows)
- o-e-sort blocks



P A R C

sorting in constant time



Sort blocks:

broadcast (1)

rank (2)

broadcast (1)

Complete sort:

sort blocks

all-to-all (2)

sort blocks

all-to-all (2)

o-e-sort blocks

P A R C




- **better than PRAM --- but useless!!!**



P A R C

Kunde's all-to-all mapping




P A R C

vertical all-to-all




P A R C

horizontal all-to-all



P A R C

Use of bus



P Å R C

sorting in optimal time Kunde / Schröder

$(k/2)^2$ steps

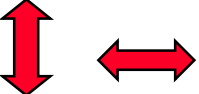
$$k = n^{1/3}$$

each step takes $n^{1/3}$ time

$$\rightarrow T = n/4$$

 x 2

 x 2

 /2

$$T_{\text{all-to-all}} = n/2$$

Sorting:

sort blocks ($O(n^{2/3})$)

all-to-all ($n/2$)

sort blocks ($O(n^{2/3})$)

all-to-all ($n/2$)


sort blocks ($O(n^{2/3})$)

time: $n + o(n)$

P A R C

Why optimal?

Sorter for n keys



Bisection of data with k wires



Sorting time $> n/k$



Use of theorem

1.) n keys on a $k \times k$ RM:

Time $\geq n/k$

Proof:

**Wherever the data is stored there is always a bisection of length k
-- this can be demonstrated sweeping left right through the array.**

Q.e.d.

2.) $n \times n$ keys on an $n \times n$ RM:


Time $\geq n$.

Proof: trivial

P Å R C


$n + o(n)$

Optimal --- but ...



P A R C

enumeration / prefix sum



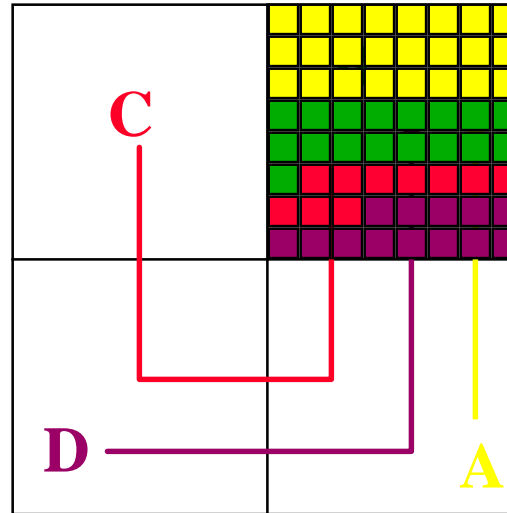

time: $O(\log n)$

wire efficiency ! -- (compared with tree)
1/2 number of processors

P A R C

ABCD-routing


- move and smooth




Row-major enumeration of A, B, C and D packets within each quadrant in time $4 \log n$. Determine destination position of each packet.

P Å R C

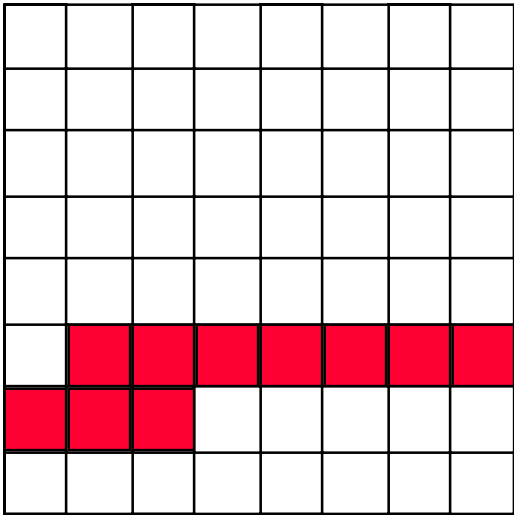
elementary steps




move



collect

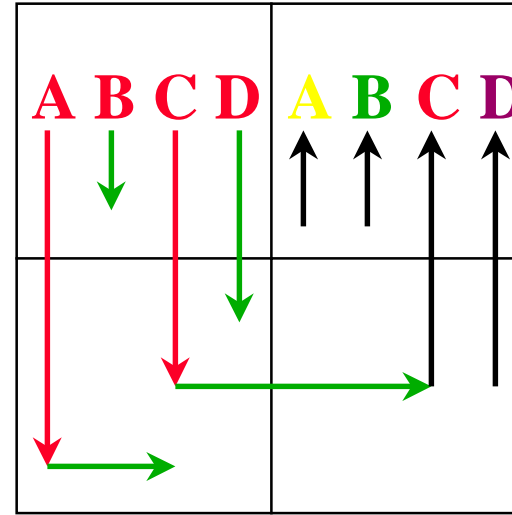



smooth



P A R C

time analysis




time: $3 \times n/2$

$$T=3n+o(n)$$

P A R C

$$T < 2n$$



4 destination squares

time: $3n + 4 \log n$

16 destination squares

time: $2n + 16 \log n$

64 destination squares

time: $12/7 n + 64 \log n$

mesh-diameter: $2n$



enough of routing/sorting

Constant factor !

Can we do better ?

What kind of problems ?

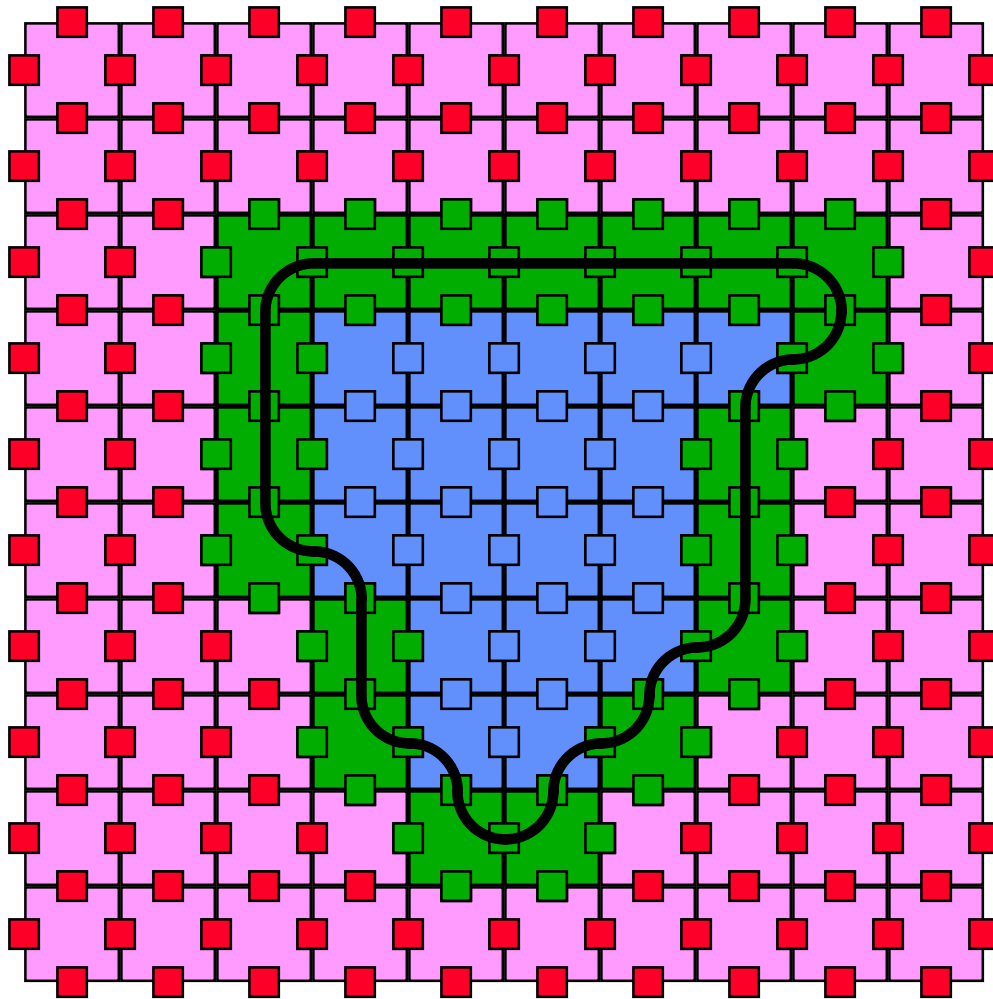
Image processing

Sparse problems !

- **Border following**
- **Edge detection**
- **Component labeling**
- **Skeletons**
- **Transforms**

P Å R C

Component labelling



Object
Define border
(candidates)
Set bus

While own label is
not received:

- 1.) **Candidates** brake bus
and send their label
 - a) clockwise
 - b) anti-clockwise
- 2.) **Candidates** switch off
and restore bus if they
see smaller label


Time: $O(1)$ -- $O(\log n)$

- Wavelet transform: Time $\log n$ on RM
-- time n on mesh
- FFT: Time n on RM and mesh
- Hough transform: Time $m \times \log n$ on RM
-- time $m \times n$ on mesh

systolic matrix multiplication


P A R C

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$



P A R C

sparse matrix multiplication



Time: n ($n \times n$ mesh)

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$

A and B column sparse (k^2)

A and B row sparse (k^2)

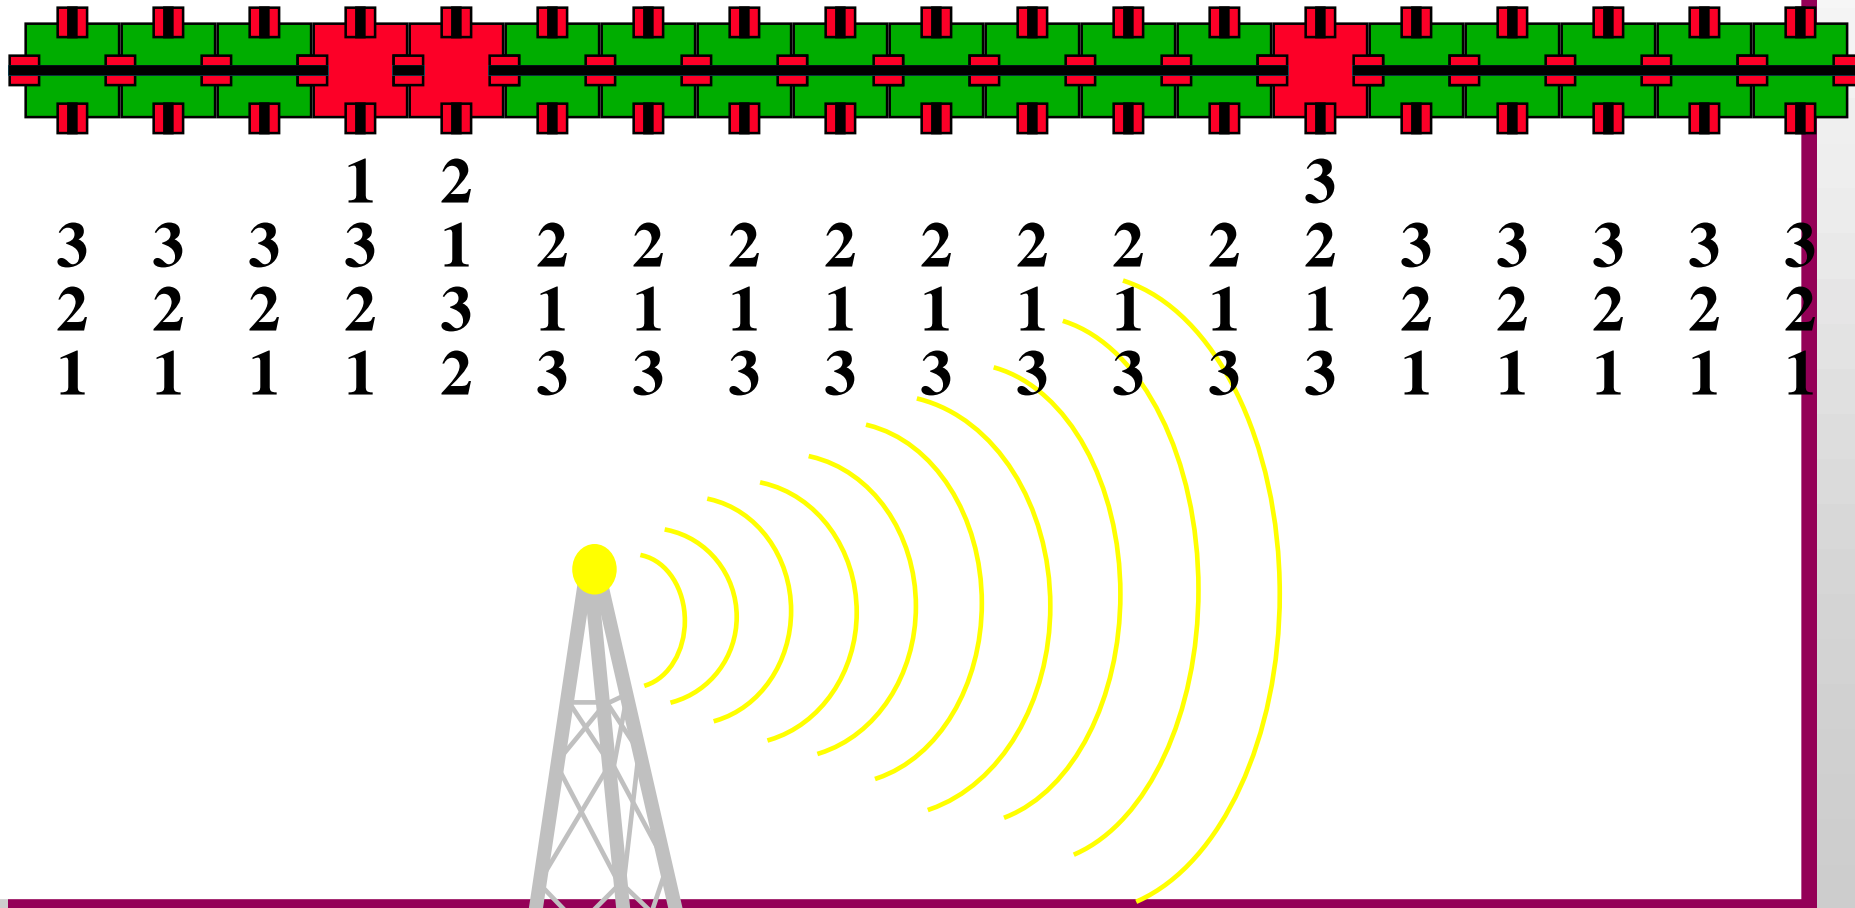
A row sparse, B column sparse (k^2)

A column sparse, B row sparse ($k\sqrt{n}$)

PARC

unlimited bus length

- ring broadcast



P A R C

A row-sparse B column-sparse

Repeat k times

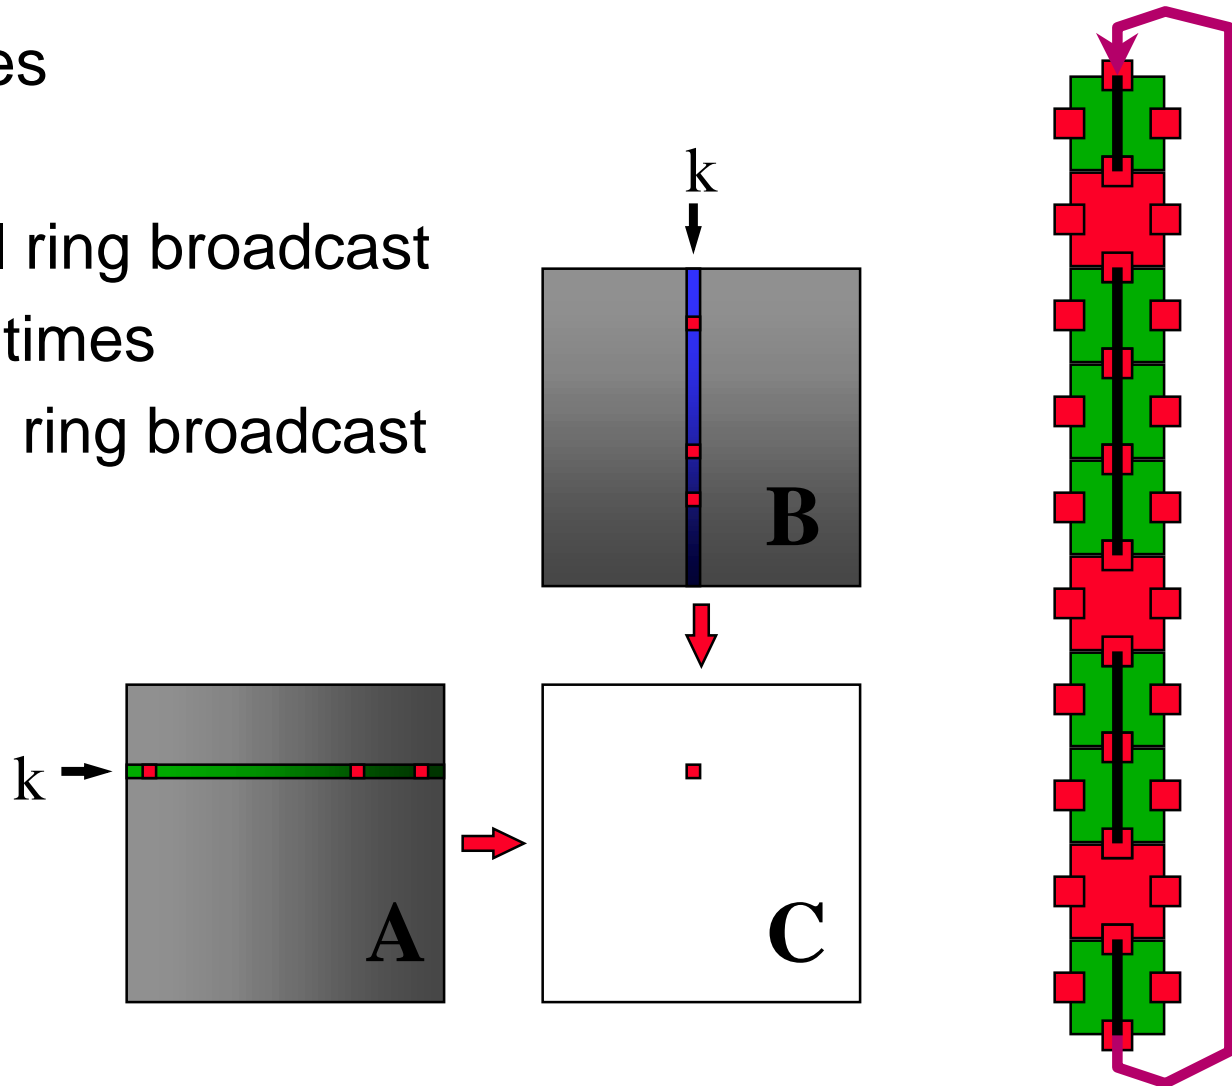
Begin

horizontal ring broadcast

Repeat k times

vertical ring broadcast

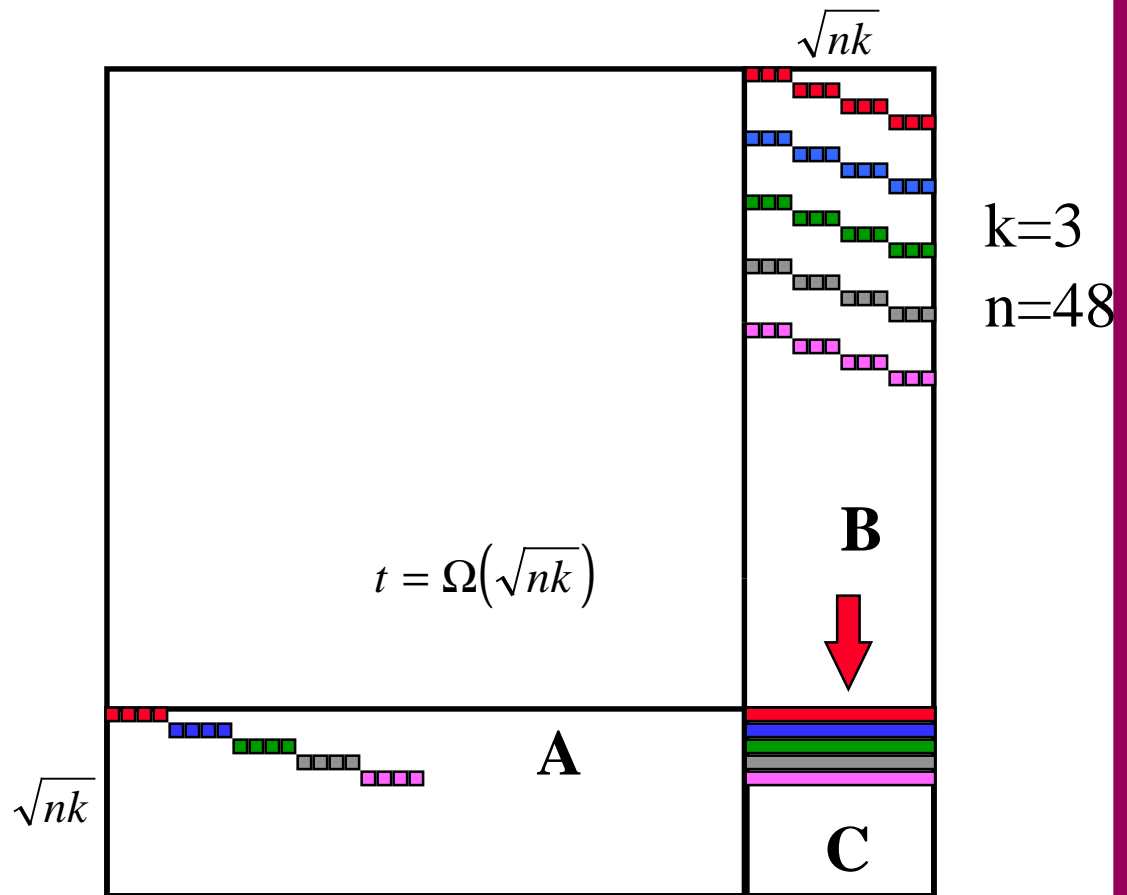
End.



P A R C

lower bound (c,r)

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$



P A R C

splitting the problem

Repeat k times

Begin

vertical ring broadcast

Repeat s times


horizontal ring broadcast

End.


$$A = A^s + A^r$$

$$C = A^s B + A^r B$$

$$C = C^s + C^r$$

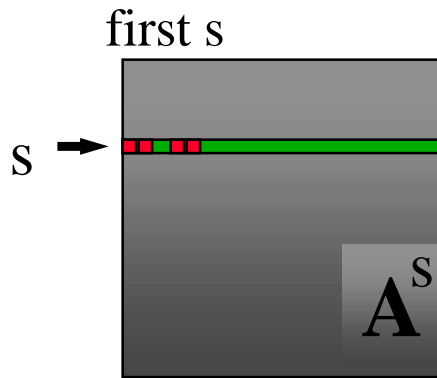


time: ks



P A R C

CR



$$A = A^s + A^r$$

A has nk non-zero elements \rightarrow

A^r has at most nk/s non-zero rows \rightarrow

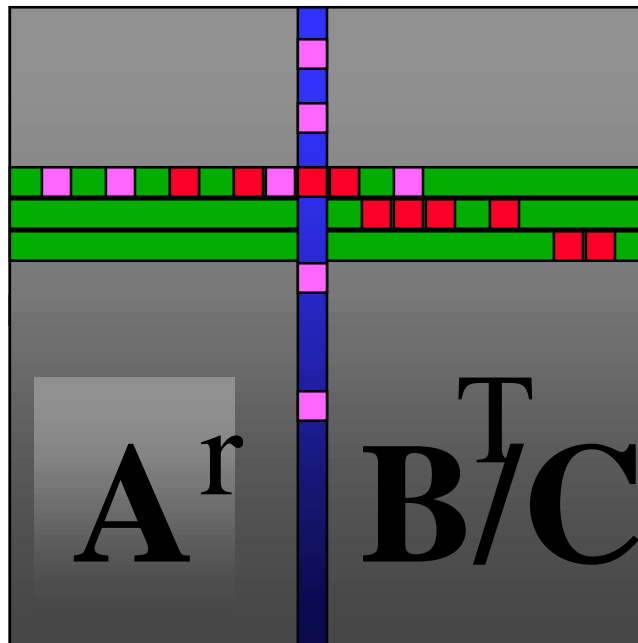
for $s = \sqrt{n}$ A^r has at most $k \sqrt{n}$ non-zero rows.

$A^s B$ is a CC- problem \rightarrow it takes time $k \sqrt{n}$.

P A R C

$A^r B$ calculating products

k B-elements



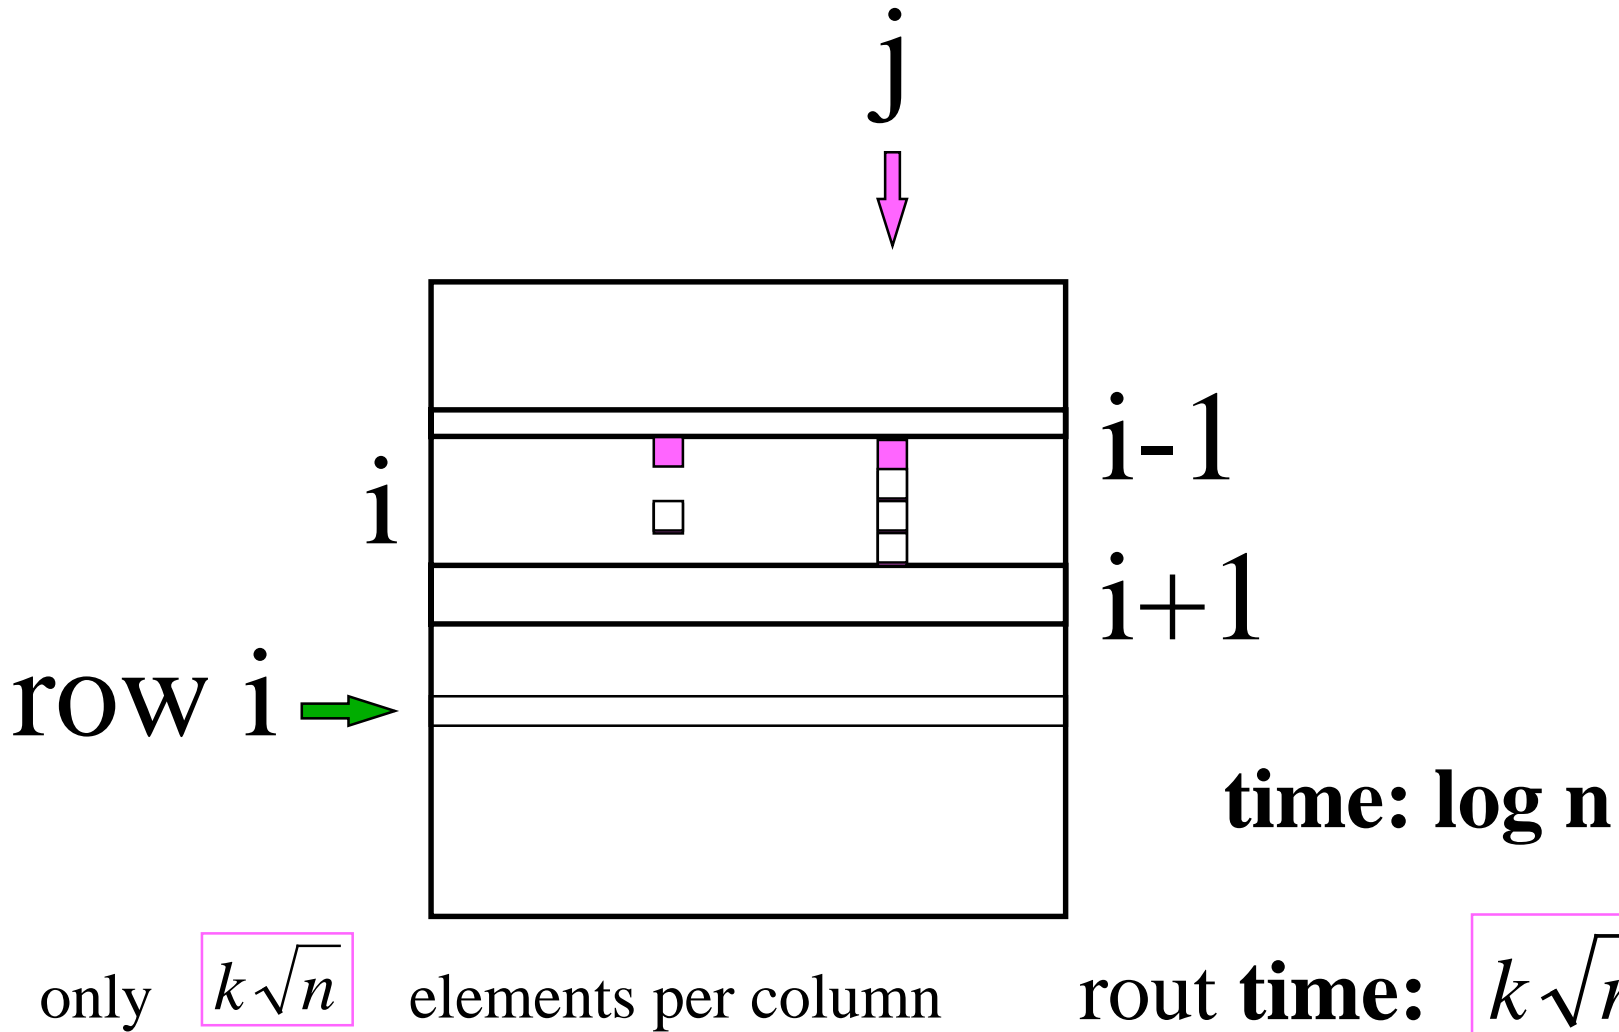
k A-elements

time: k^2

$k^2 n$ elements


P Å R C

column sum



P A R C

routing within columns



route **time:**

$$k\sqrt{n}$$

P A R C

Reconfigurable architectures

Reconfigurable mesh ?

constant diameter !




Physical laws!

P A R C

Physical limits

$c=300\,000\text{ km/sec}$




- 30cm/ns
- on chip: 1cm/ns
- --> **bounded bus length**



P A R C

creating main stations



time: k

P A R C

A row-sparse B column-sparse

Create main stations $1, \dots, k$ for A and B (time: $n/l+k$)

For $i=1, \dots, k$ do


Begin

horizontal ring broadcast i of A

For $j=1, \dots, k$ do

vertical ring broadcast j of B

End.



P A R C

A and B column-sparse

Create main stations $1, \dots, k$ for A (time: $n/l+k$)

For $i=1, \dots, k$ do


Begin

horizontal ring broadcast i

k bounded vertical broadcasts of products


merging new products

End.



P Å R C

remove minor stations



			1	2										3					
				1	2	2									3	3	3		
3	3				1	1	2	2	2	2	2							3	3
		3	3	3	3	3	1	1	1	1	1	2	2	2	2	2			
2	2						3	3	3	3	3	1	1	1	1	1	1	2	2
1	1	2	2	2	2	2						3	3	3	3	3	3	1	1

PARC

results

Time: n ($n \times n$ mesh)

A and B column sparse (k^2) ($k^2 + 2n/l$)

A and B row sparse (k^2) ($k^2 + 2n/l$)

A row sparse, B column sparse (k^2) ($k^2 + n/l$)

A column sparse, B row sparse ($3k\sqrt{n}$) ($+11n/l$)

(Kunde, Middendorf, Schmeck, Schröder, Turner)

- image processing
 - sorting
 - routing
- load balancing

better than the mesh !

(Kapoor, Kunde, Kaufmann, Schroeder, Sibeyn)

- The RM is in some cases “better” than PRAM
- The RM is always at least as “good” as mesh
- The RM is often “better” than the mesh



The End