The Future of Parallel Computing

Special Purpose Mesh Architectures

Heiko Schröder, 1998

Contents

- **Why meshes ???**
- **Application specific parallel mesh architectures**
  - Systolic Arrays
  - Instruction Systolic Arrays
  - PIPS
  - Reconfigurable mesh
  - Optical Highway

Fine grain
1983

Coarse grain
1997

Physical limits

c=300 000 km/sec

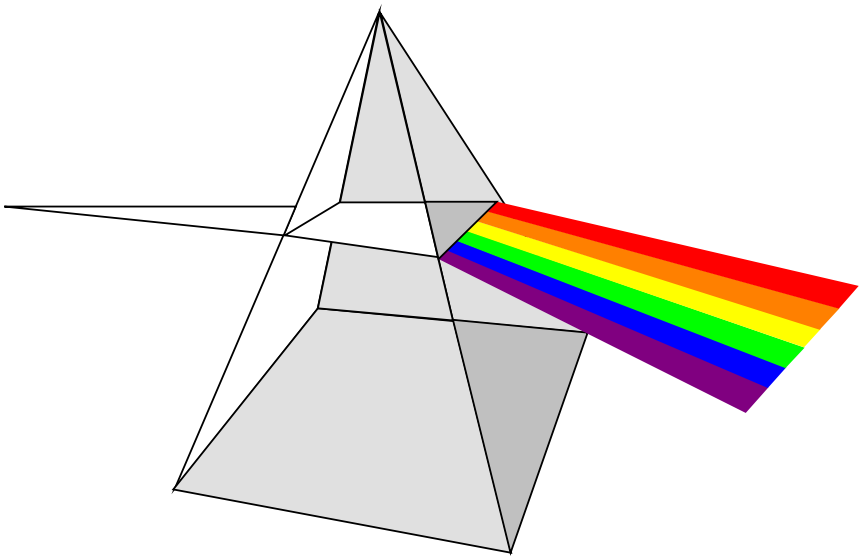

- 10^{12} OPS -- 0.3 mm/OP
- 1000 PEs with 10^{9} OPS -- 30cm/OP
- massive parallelism
- distributed memory

Processor power

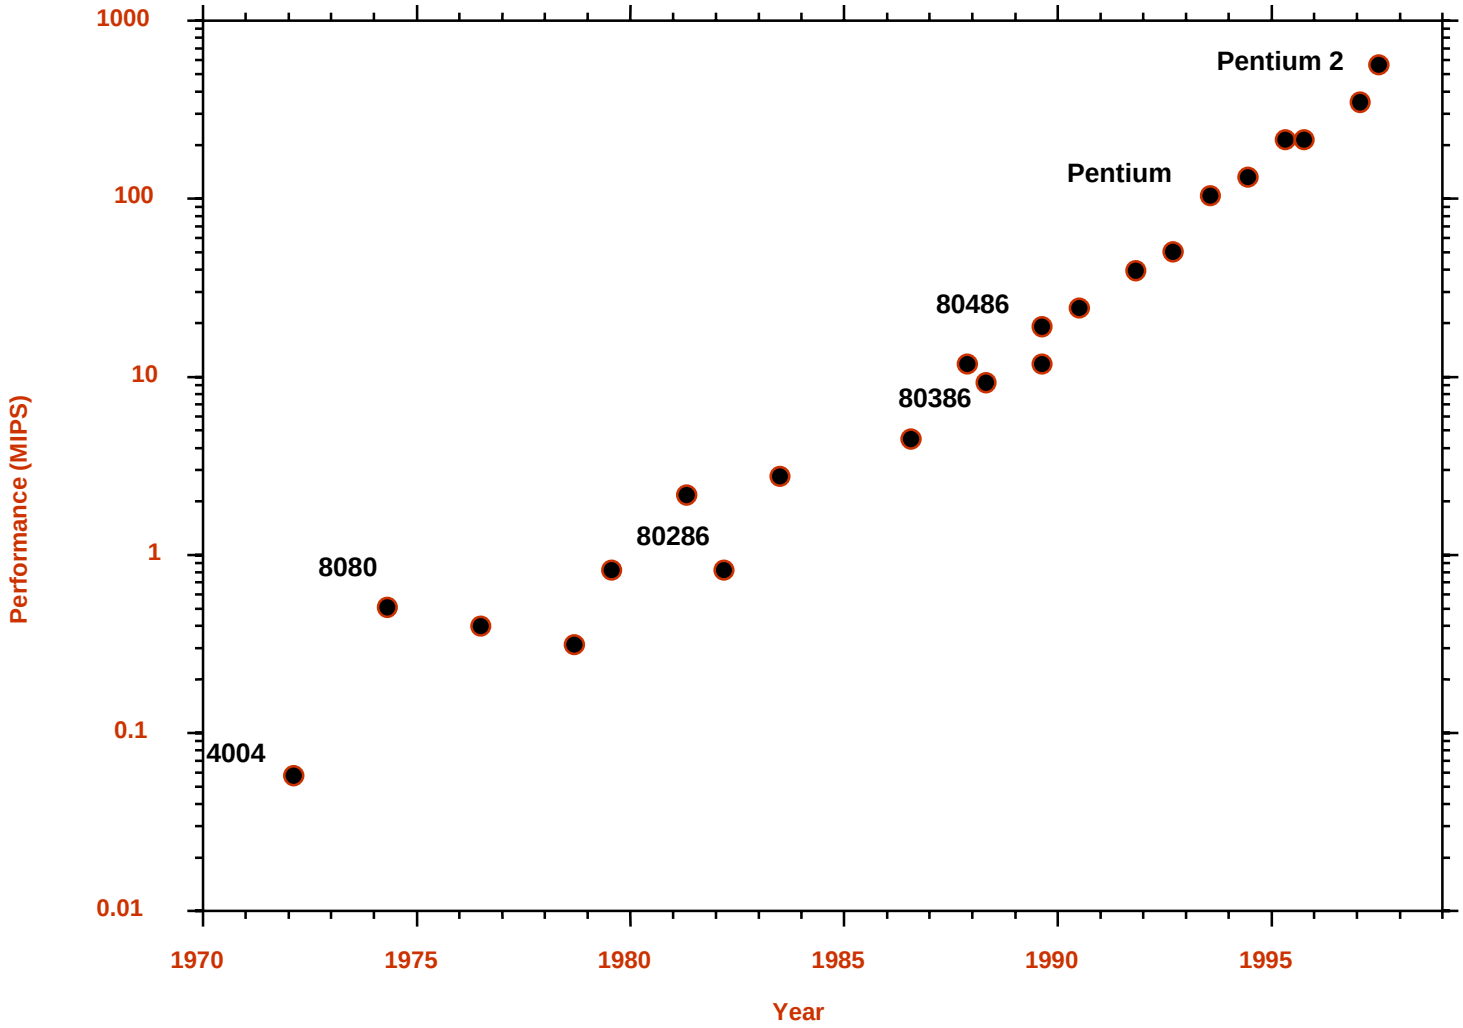

PÅRC

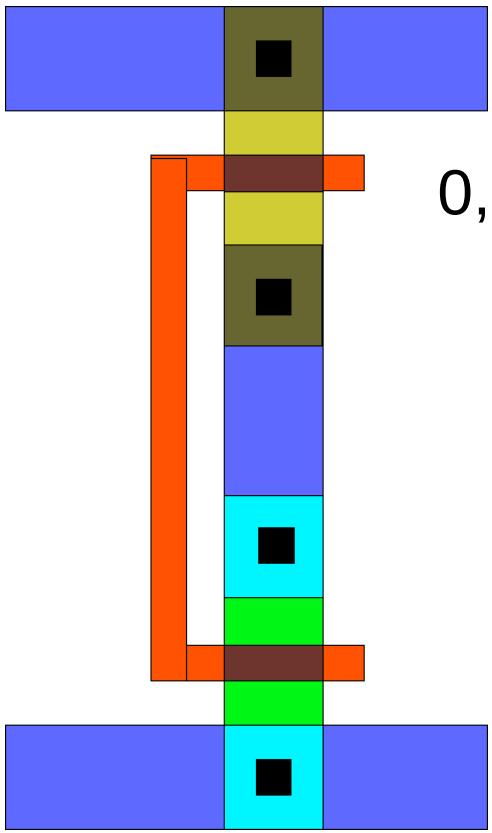

0,5 µ

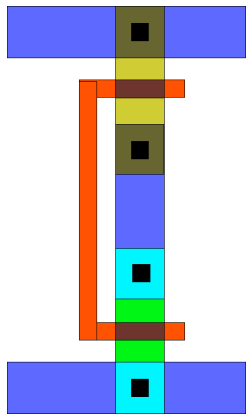

0,25 µ

Scaling
Faktor 2:

- 1/2 width
- 1/2 hight
- 1/2 switching time

8 x performance!

CMOS transistors

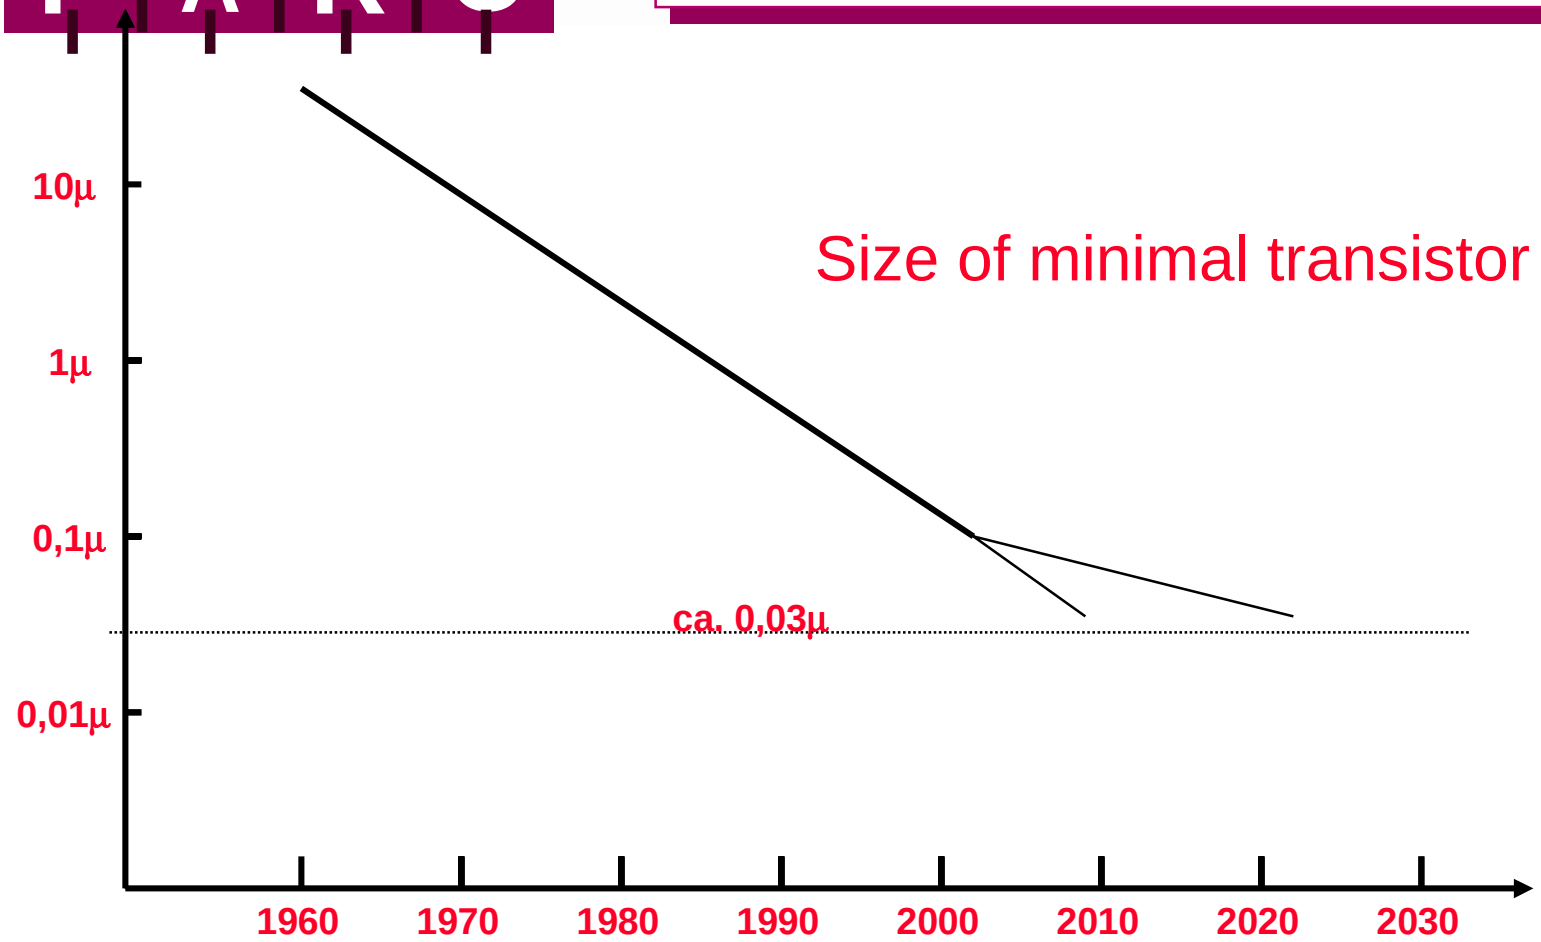

Mesh/Torus

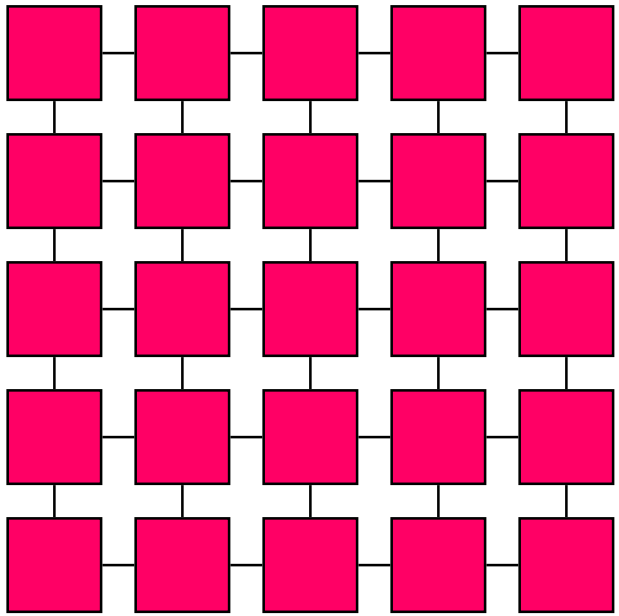

2D mesh

diameter $\sqrt{n}$

bisection width $\sqrt{n}$

Hypercube

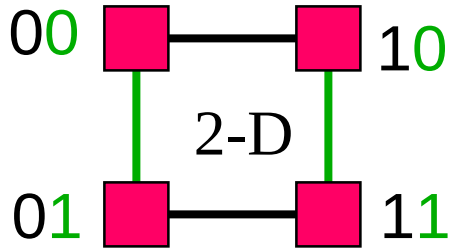

diameter log n
bisection width n

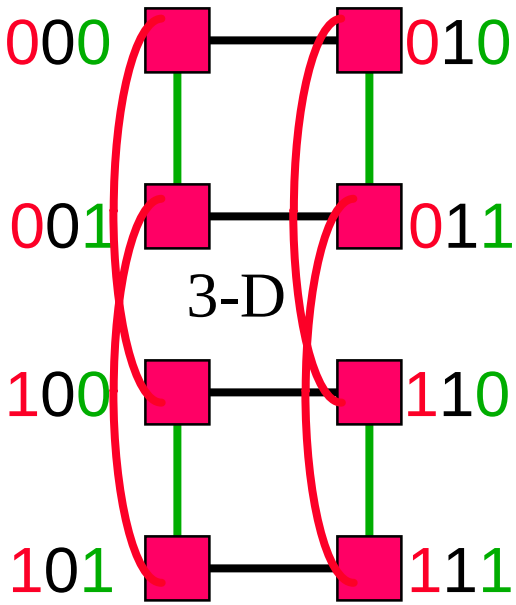

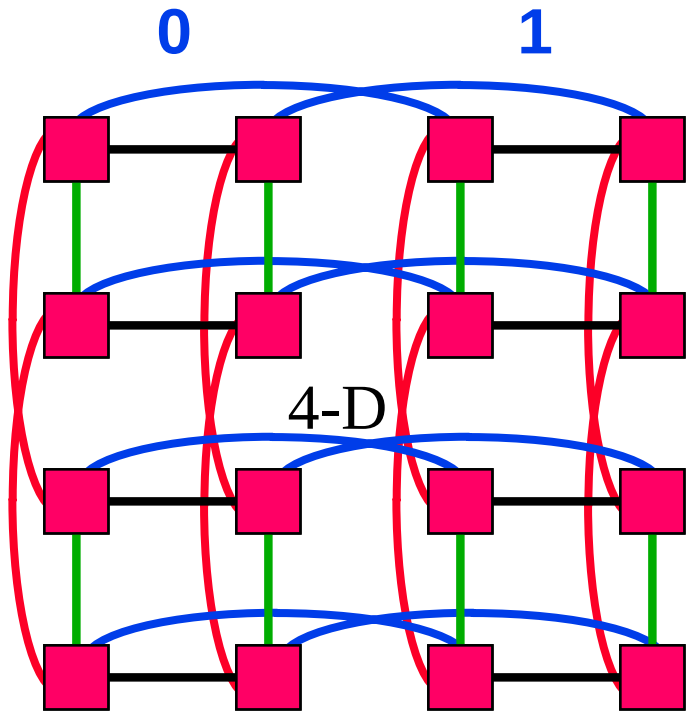

VLSI

Very

Large

Scale

Integration

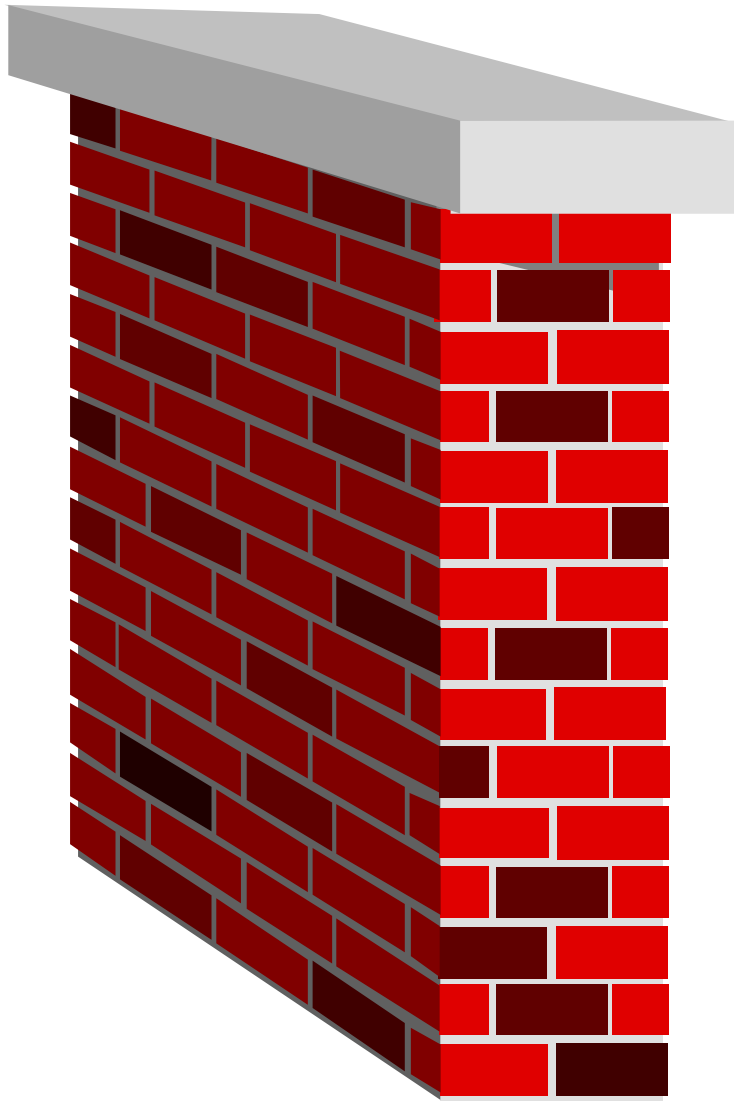

- **simple cells**
- **few types**
- **regular architecture**
- **short connections**

mesh -- torus

Pin limitations

diameter 256

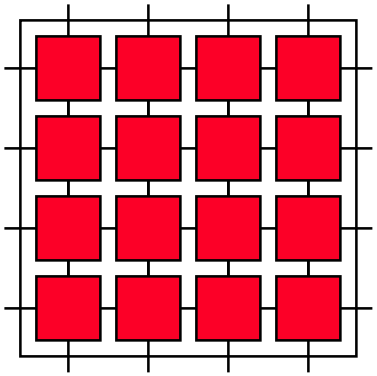

16 pins

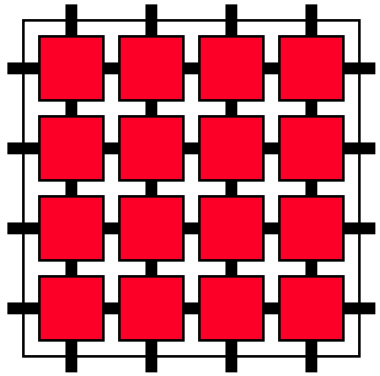

16x16 pins

diameter 16

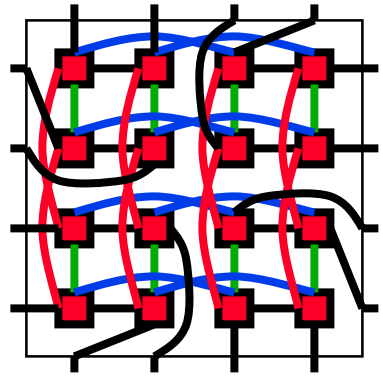

16x12 pins

Bisection width

Bisection width 256

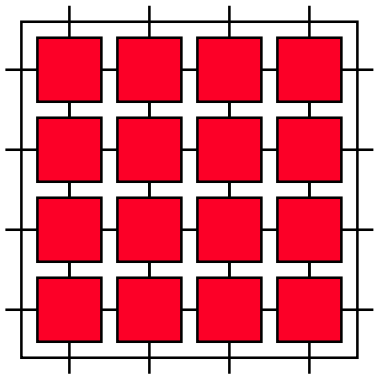

25 cm

Bisection width 32K

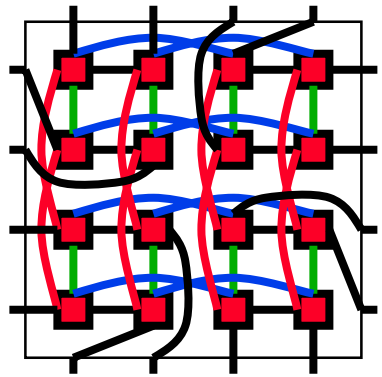

32 m

Programming

- **SA** --- **S**ystolic **A**rray
- **SIMD** --- **S**ingle **I**nstruction **M**ultiple **D**ata
- **ISA** --- **I**nstruction **S**ystolic **A**rray
- **MIMD** --- **M**ultiple **I**nstruction **M**ultiple **D**ata

parallel merge

initial situation:

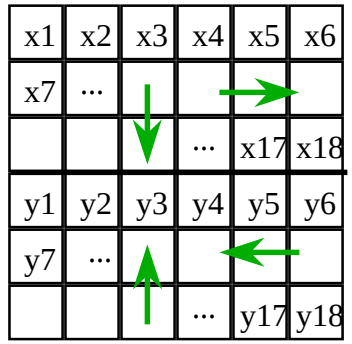

1.) sort columns

(odd-even-transposition sort)

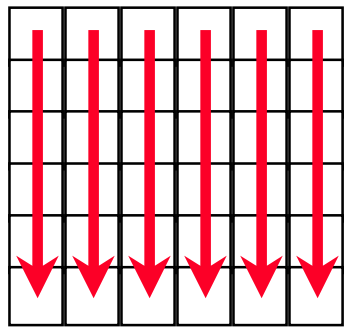

2.) sort rows

(odd-even-transposition sort)

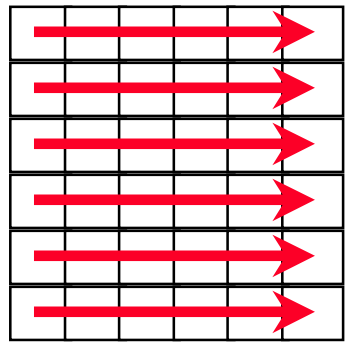

sorted !!!!

0-1 principle

- The 0-1 principle states that if all sequences of 0 and 1 are sorted properly than this is a correct sorter.
- The sorter must be based on moving data.

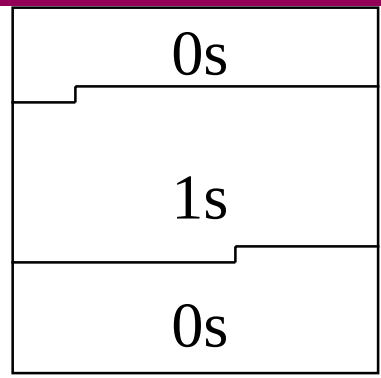

initially

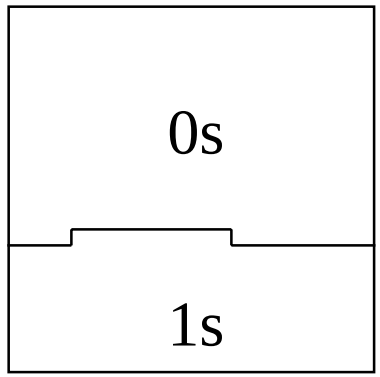

after vertical sort

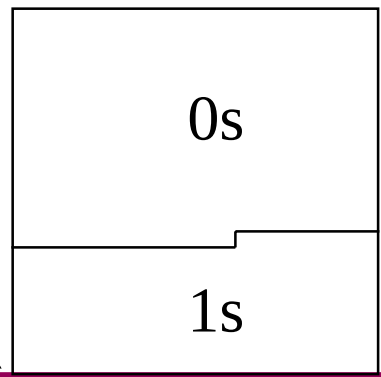

after horizontal sort

MIMD-mesh (clocked)

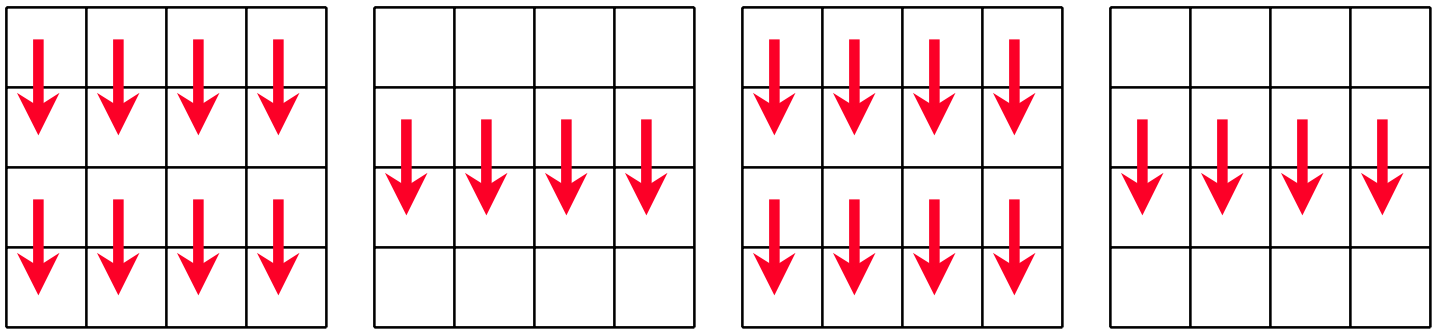

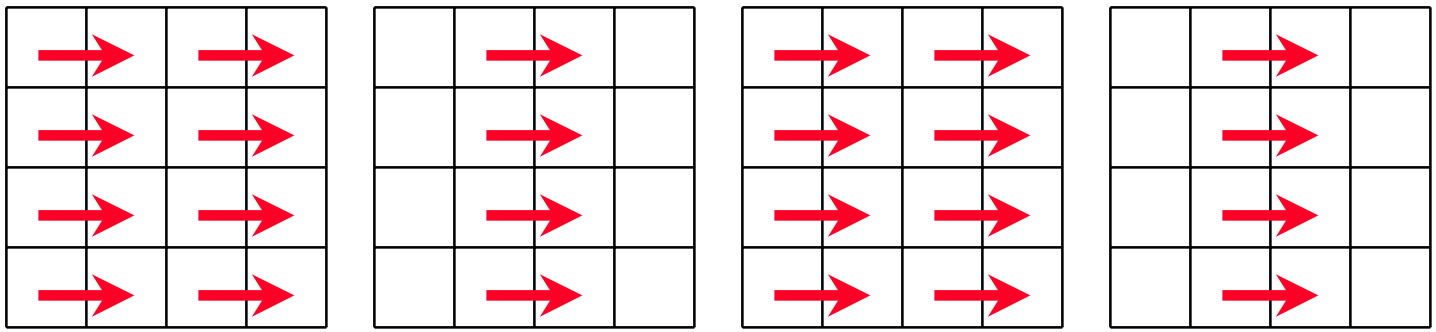

min 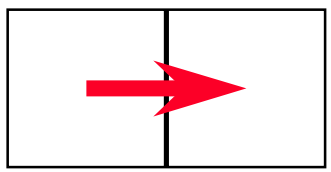 max

Time: 2n

systolic merge

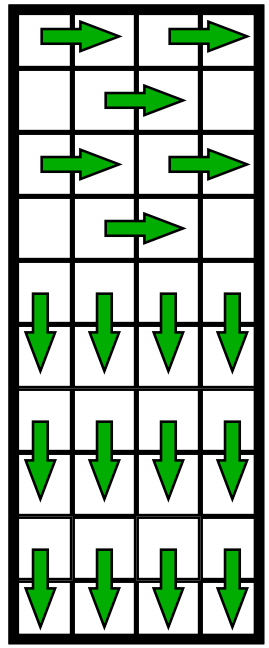

1	3	3	4
5	5	6	7
9	8	8	7
4	4	3	2

systolic merge

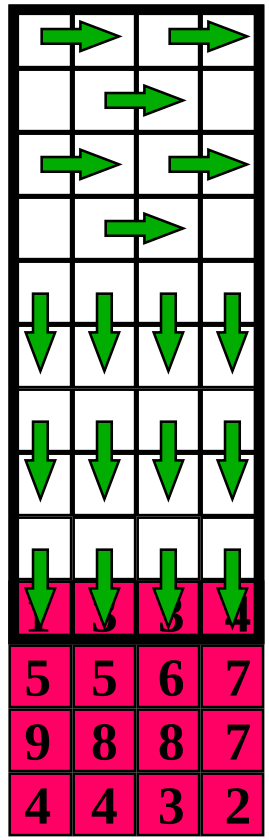

systolic merge

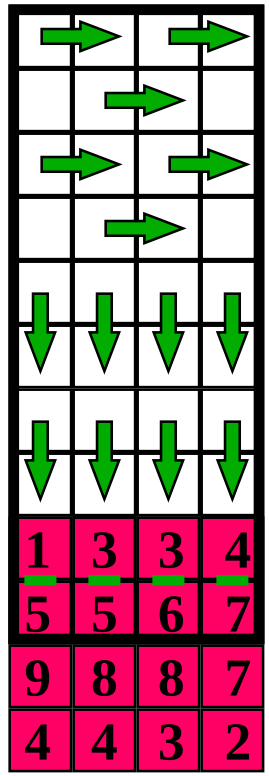

systolic merge

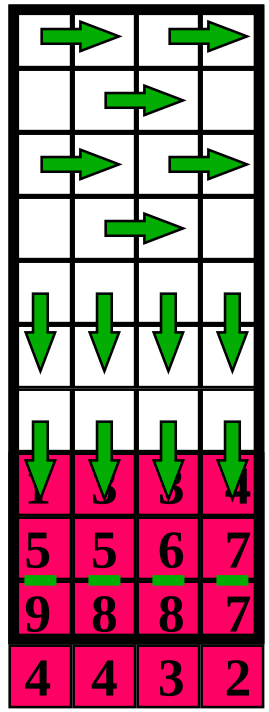

systolic merge

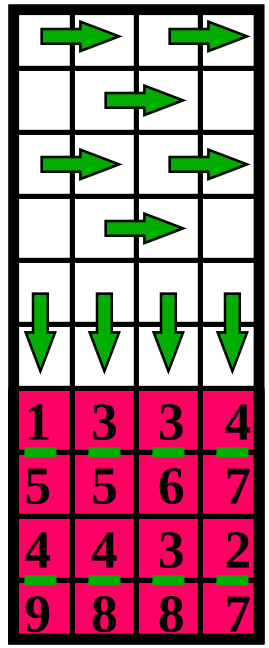

systolic merge

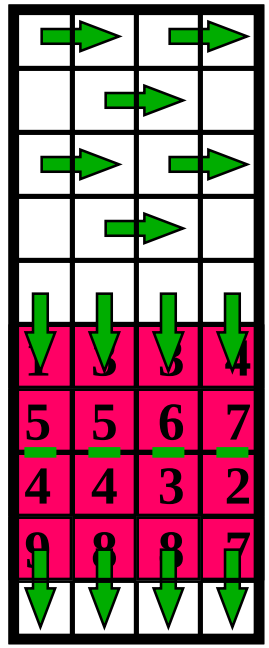

systolic merge

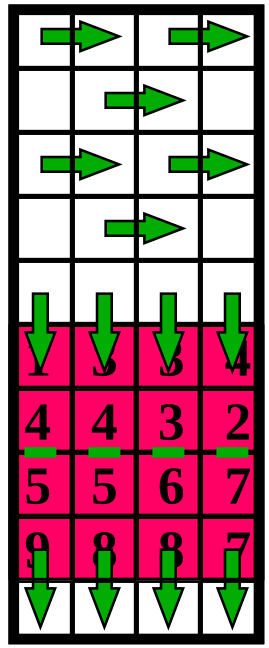

systolic merge

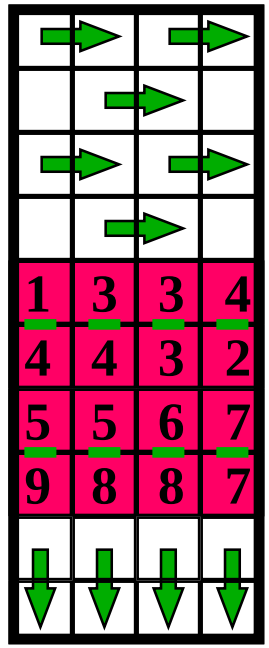

systolic merge

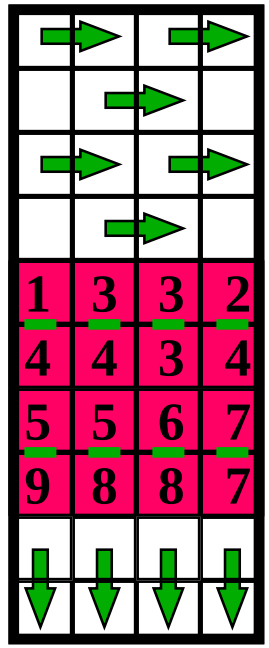

systolic merge

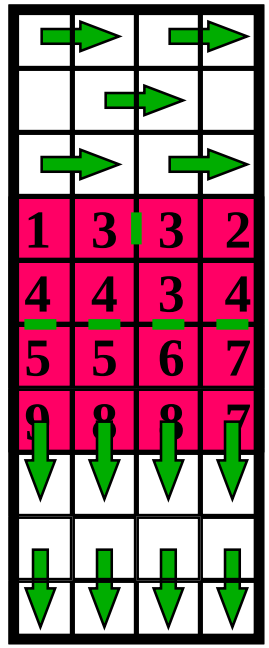

PÅRC

systolic merge

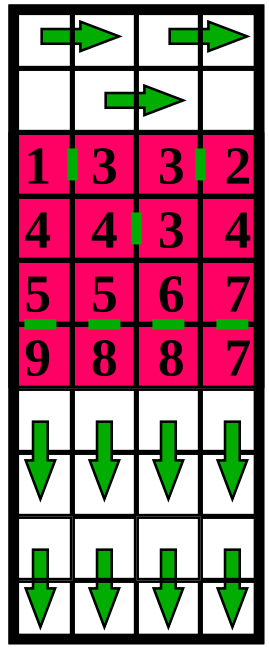

PÅRC

systolic merge

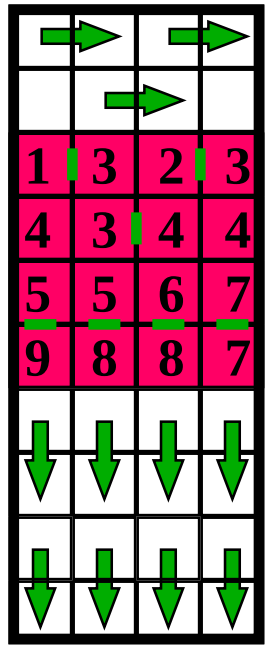

systolic merge

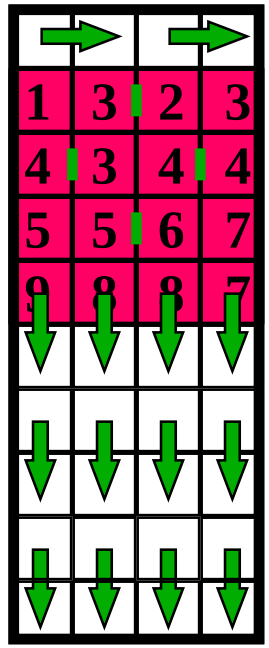

systolic merge

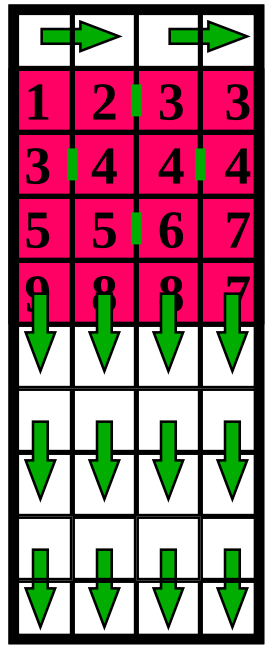

systolic merge

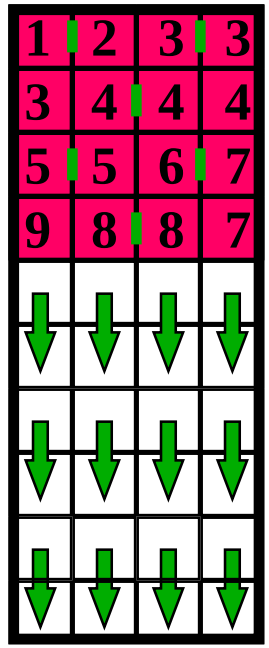

systolic merge

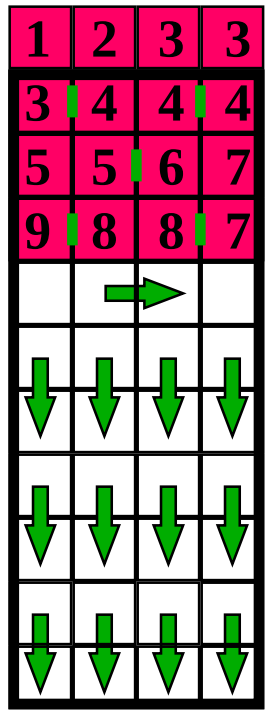

systolic merge

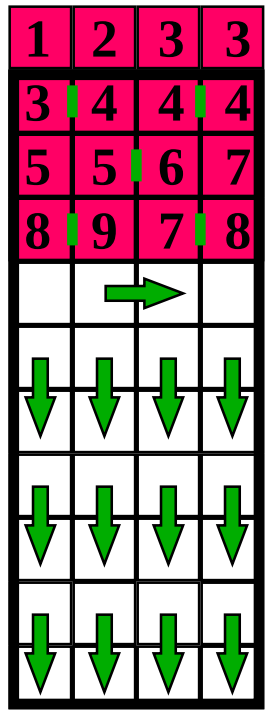

systolic merge

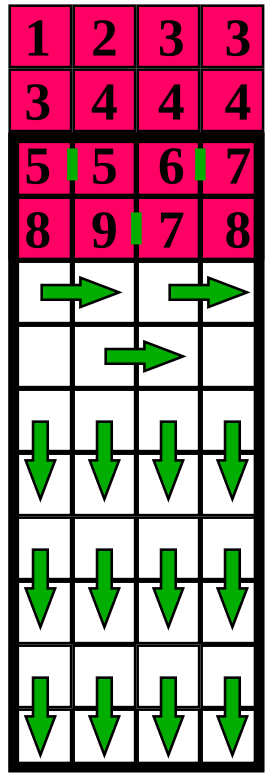

PÅRC

systolic merge

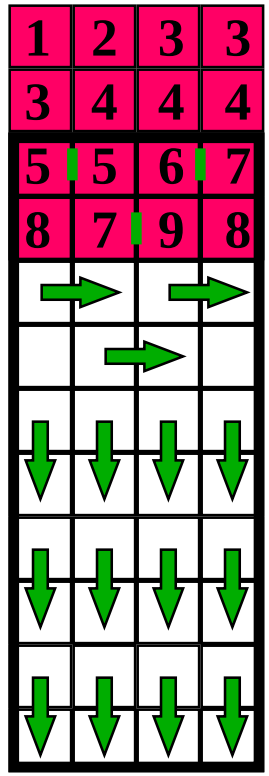

PÅRC

systolic merge

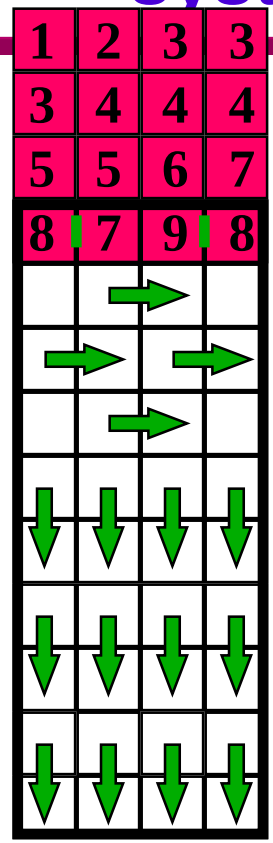

systolic merge

- **sorted !!!**

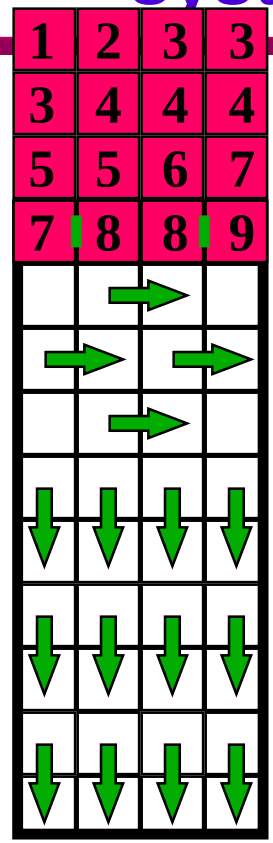

Characteristics of SAs

Extremely high cost-performance
no flexibility -- long development time

Suitable for special signal
processing tasks ???

Systolic architectures I

1. Lang, H.W., Schimmler, M., Schmeck, H., Schröder, H., "A Fast Sorting Algorithm for VLSI", Proc. 10th ICALP, Barcelona, July 1983, *Lecture Notes in Computer Science*, 154, pp408--419, 1983
2. Lang, H.W., Schimmler, H., Schröder, H., "Pattern Matching in Binary Trees on a Mesh-Connected Processor Array", *VLSI: Algorithms and Architectures*, Bertolazzi and Luccio (eds.), North Holland, pp113--124, 1984
3. Schmeck, H., Schröder, H., "Dictionary Machines for Different Models of VLSI", *IEEE Transactions on Computers*, C-34, pp472--475, 1985
4. Lang, H.W., Schimmler, M., Schmeck, H., Schröder, H., "Realistic Comparisons of Sorting Algorithms for VLSI", *Foundations of Data Organisation*, Ghosh, Kambayashi and Tanaka (eds.), Plenum Press, pp309--318, 1987
5. Schröder, H., "VLSI-Sorting Evaluated under the Linear Model", *Journal of Complexity* **4**, pp 330-355, December 1988
6. Schmeck, H., Schröder, H., Starke, C., "Systolic s^2-Way Merge Sort is Nearly Optimal", in *IEEE Transactions on Computers*, Vol 38, Nr 7, pp 1052-1056, 1989
7. Murthy, V.K., Schröder, H., "Systolic Arrays for Parallel G-Inversion and Finding Petri Net Invariants", in *Parallel Computing*, Vol. 11, Nr. 3, pp 349-359, 1989
8. E.V. Krishnamurthy, M. Kunde, M. Schimmler, H. Schröder, "Systolic algorithm for tensor products of matrices --- implementation and applications", *Parallel Computing* 13, 301-308, 1989

Systolic architectures II

9. Schimmler, M., Schröder, H., “A Simple Systolic Method to Find all Bridges on an undirected Graph”, *Parallel Computing* 12, 107-111, 1989
10. P. Lenders, H. Schröder, “A programmable systolic device for Image Processing --- based on Mathematical Morphology”, *Parallel Computing*, 13, pp 337-344, 1990
11. Schröder, H., Krishnamurthy, E.V., “Systolic Algorithms for Polynomial Interpolation and Related Problems”, *Parallel Computing*, 17, pp 493-504, 1991
12. Schröder, H., Krishnamurthy, E.V., “Systolic Algorithms for Multivariate Approximation Using Tensor Products of Basis Functions”, *Parallel Computing*, 17, pp 483-492, 1991
13. B. Beresford-Smith, J. Breckling, H. Schröder, “Systolic Codebook Generation”, *Transactions on Acoustics, Speech, & Signal Processing*, pp 144-149, Vol.1, Nr. 2, April 1993
14. Schröder, H., “Partition Sorts for VLSI”, Proceedings GI-13, Jahrestagung, Hamburg, October 1983, *Informatik-Fachberichte 73*, pp101--116, 1983
15. Schröder, H., “VLSI-gerechte Sortierverfahren”, PARS-Workshop, Erlangen, April 1984, Mitteilungen GI-PARS Nr. 2, pp134--143, 1984
16. Schröder, H. “VLSI-Sorting under the Linear Model”, Proceedings of the Tenth Australian Computer Science Conference, pp330--340, Deakin University, February 1987
17. B. Beresford-Smith, J. Breckling, H. Schröder, “Systolic Devices for Speech Processing”, CompEuro 89, Hamburg, Germany, May 1989

ISA merge

C:=min{C, CE}

C:=max{C, CW}

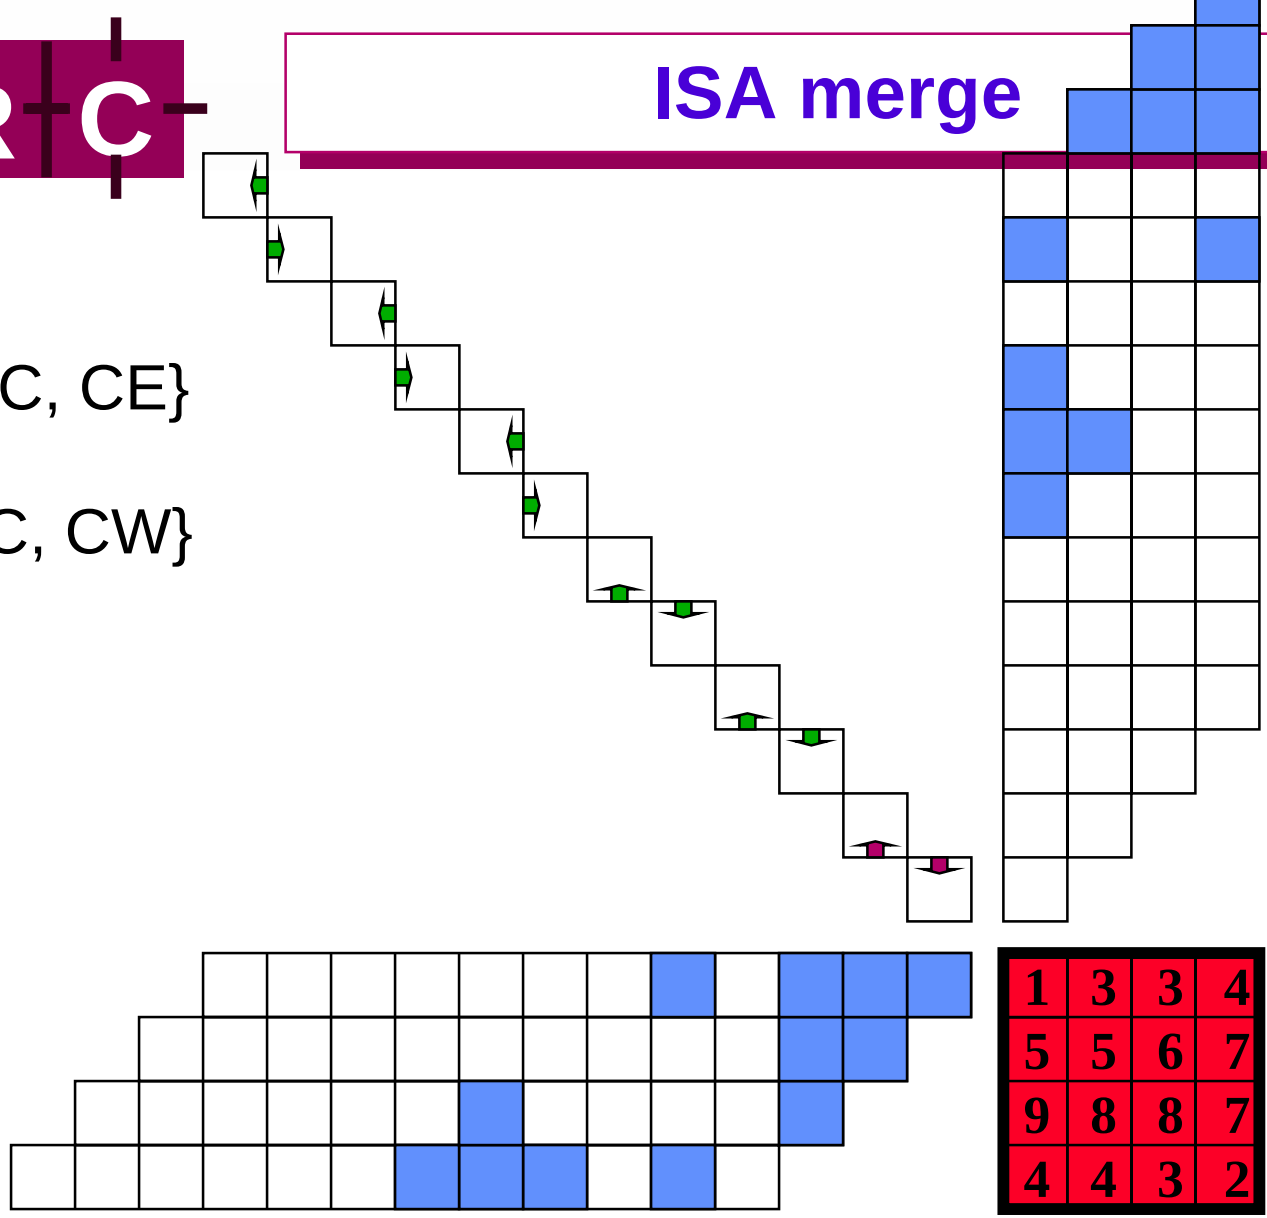

1	3	3	4
5	5	6	7
9	8	8	7
4	4	3	2

ISA merge

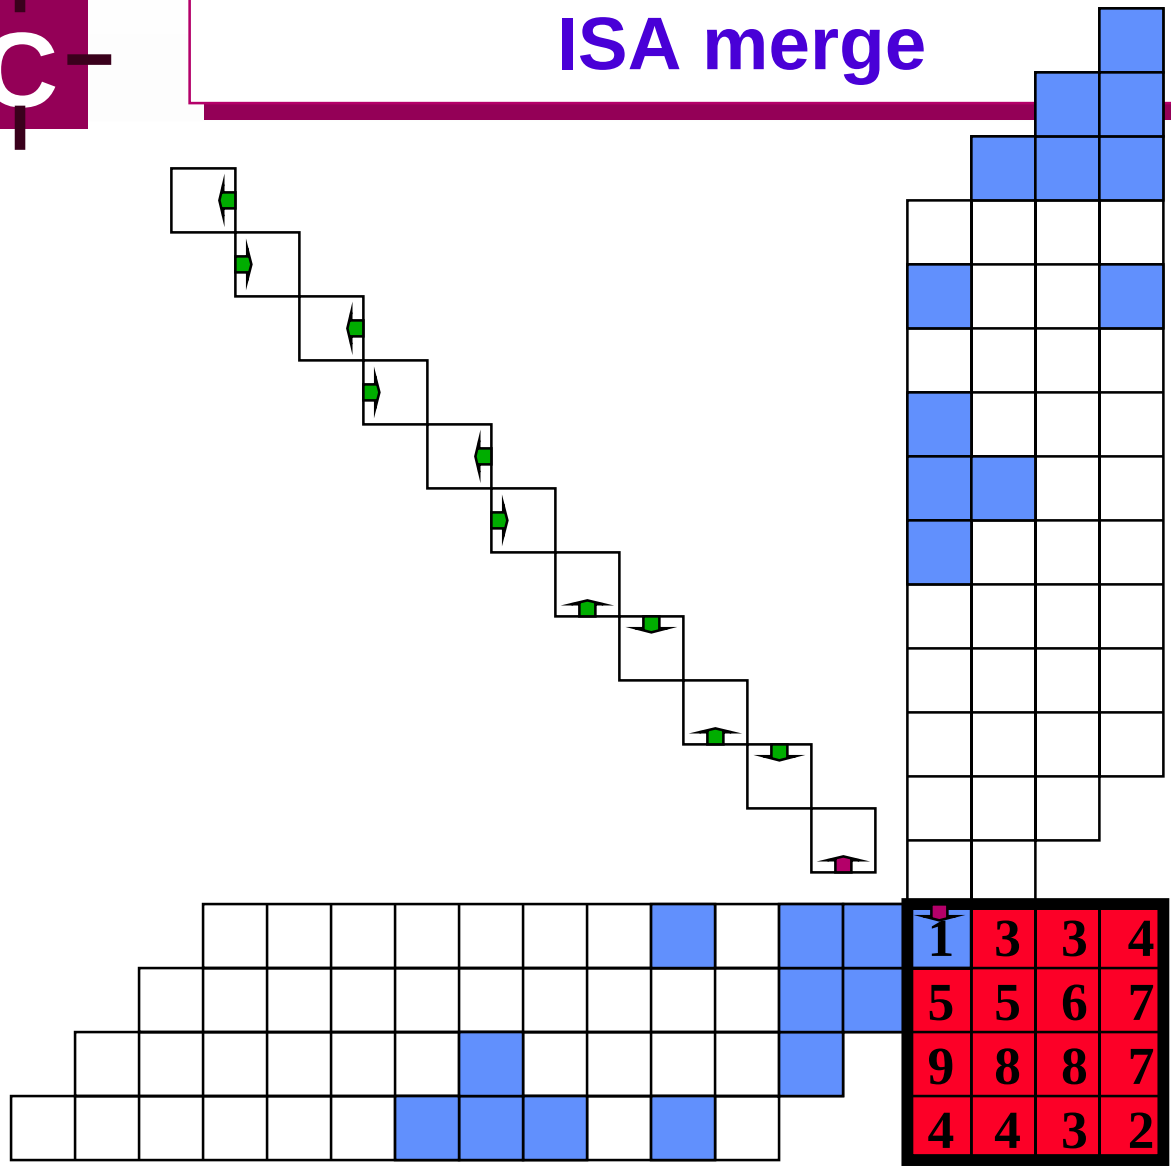

PÅRC

ISA merge

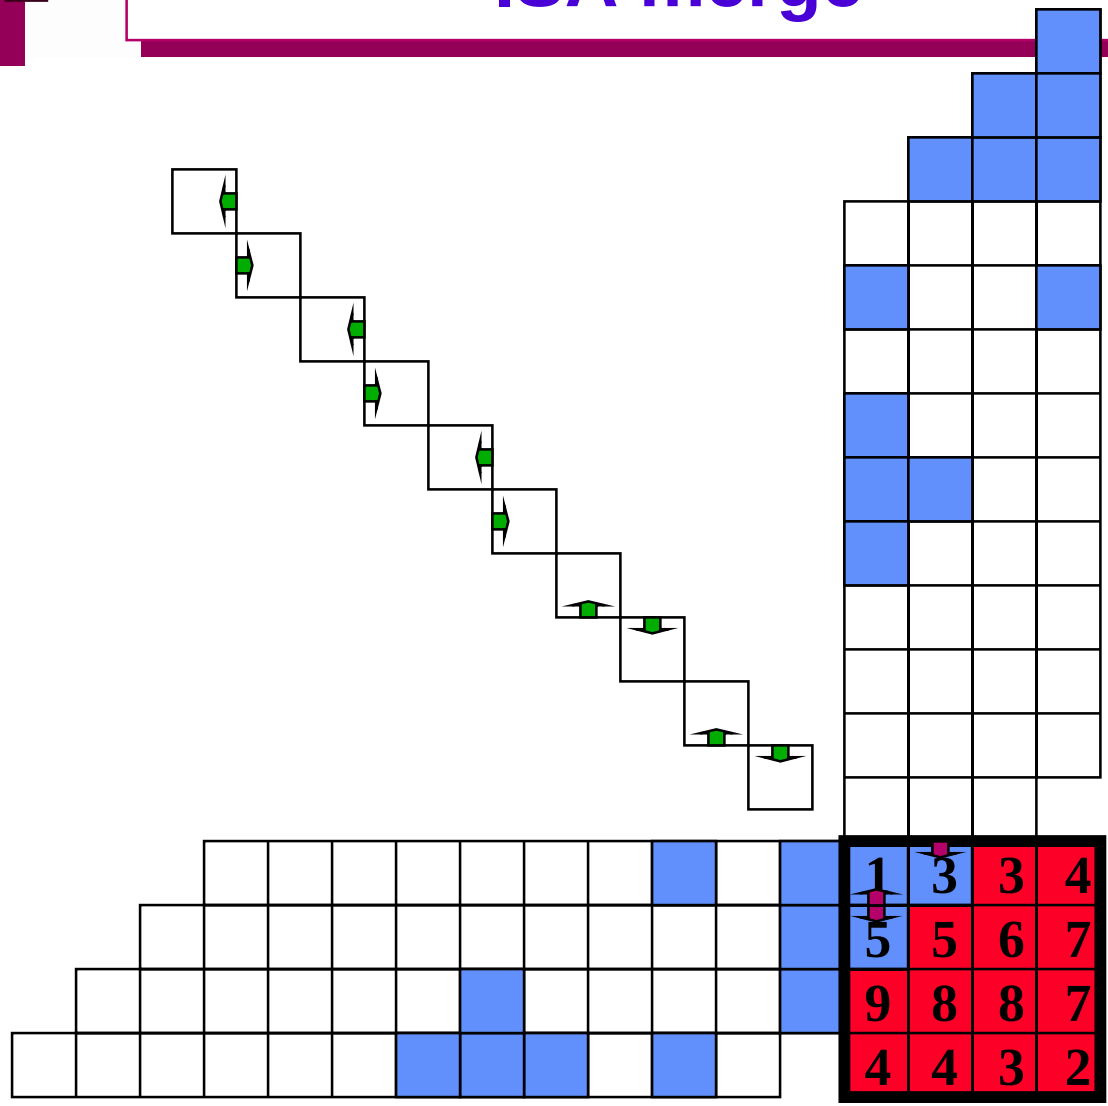

ISA merge

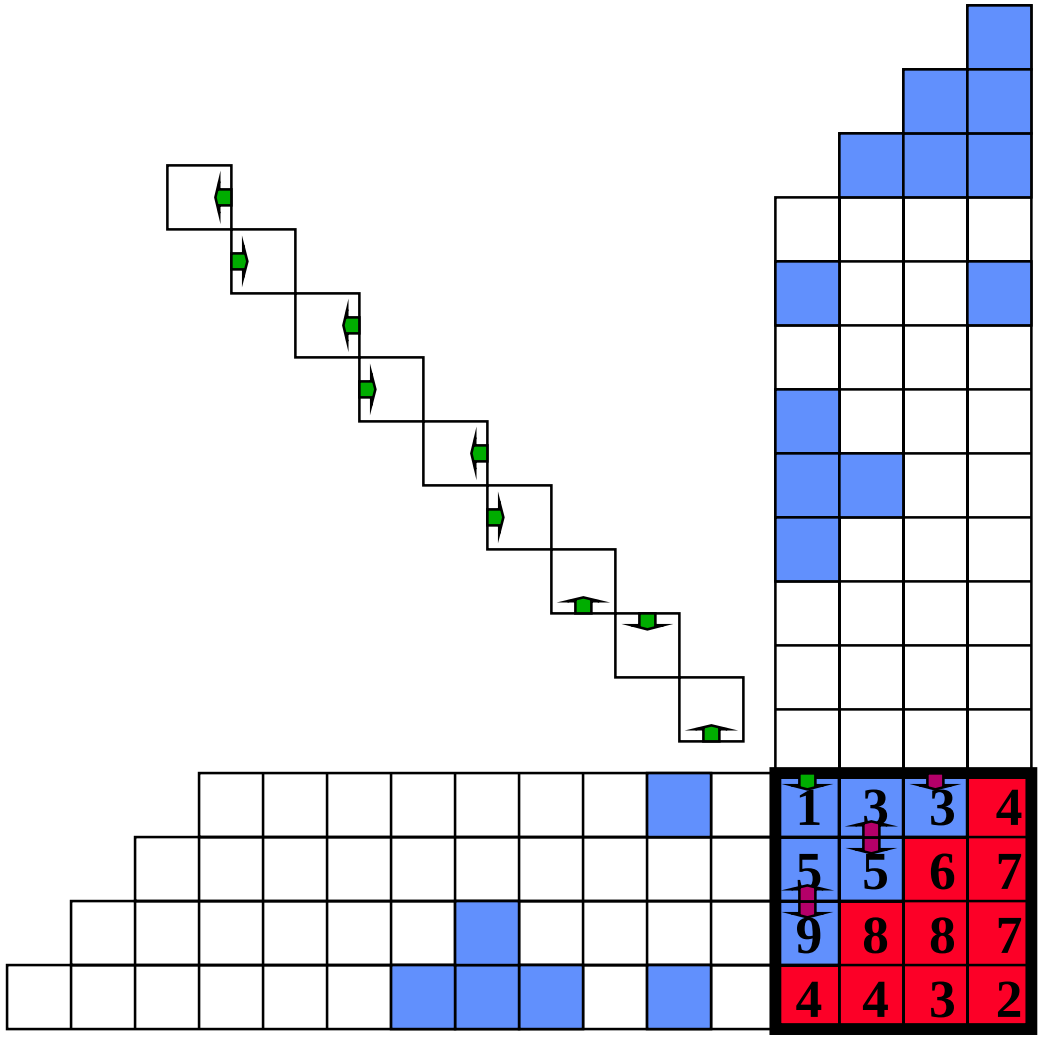

ISA merge

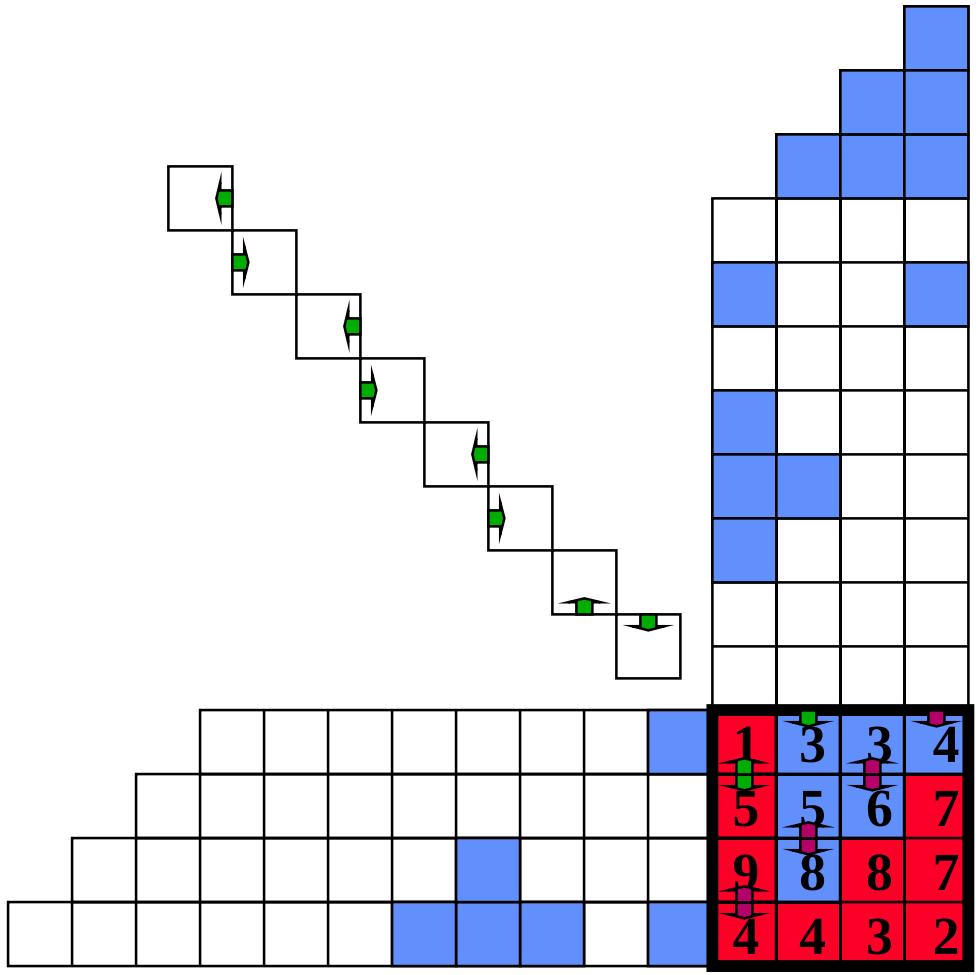

ISA merge

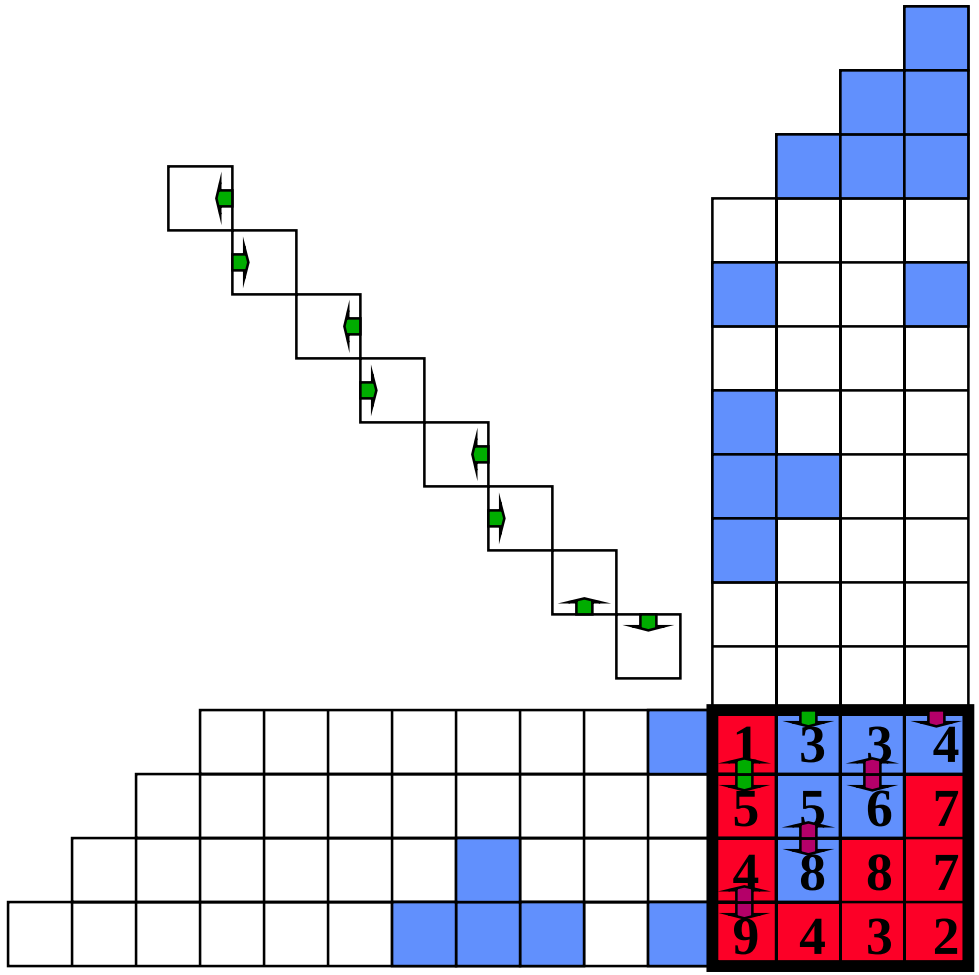

ISA merge

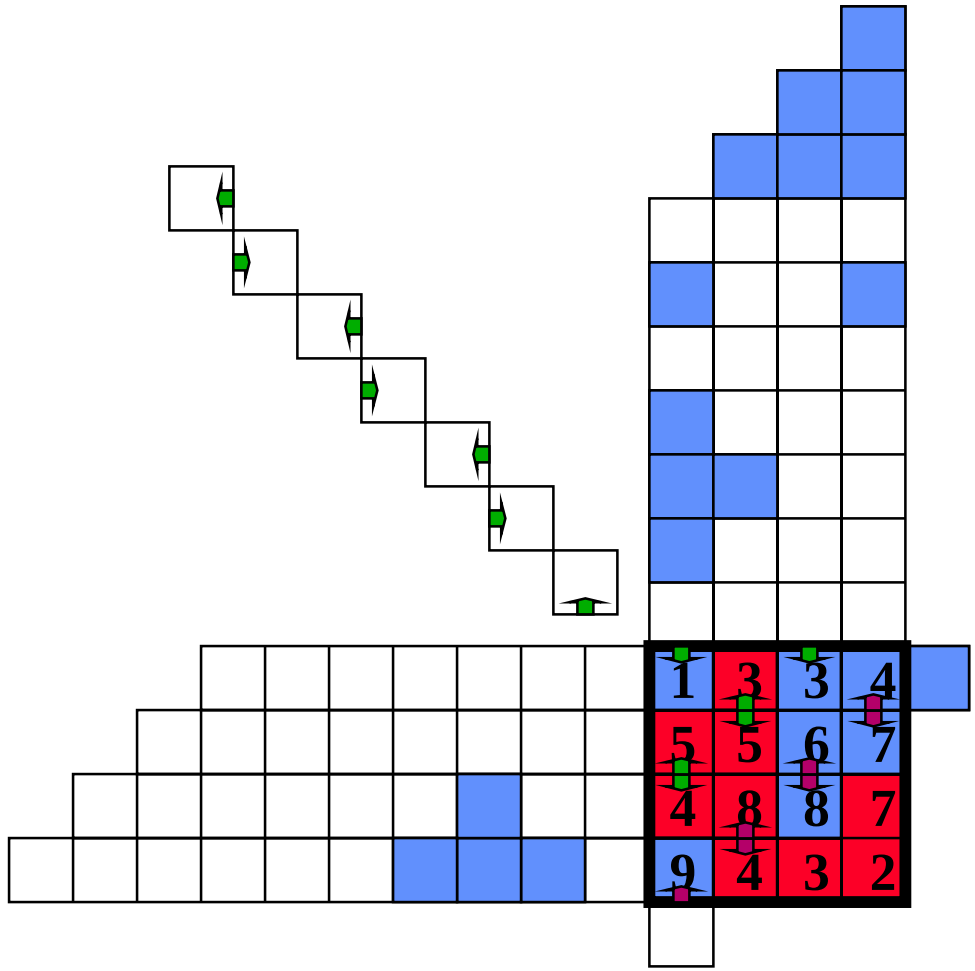

ISA merge

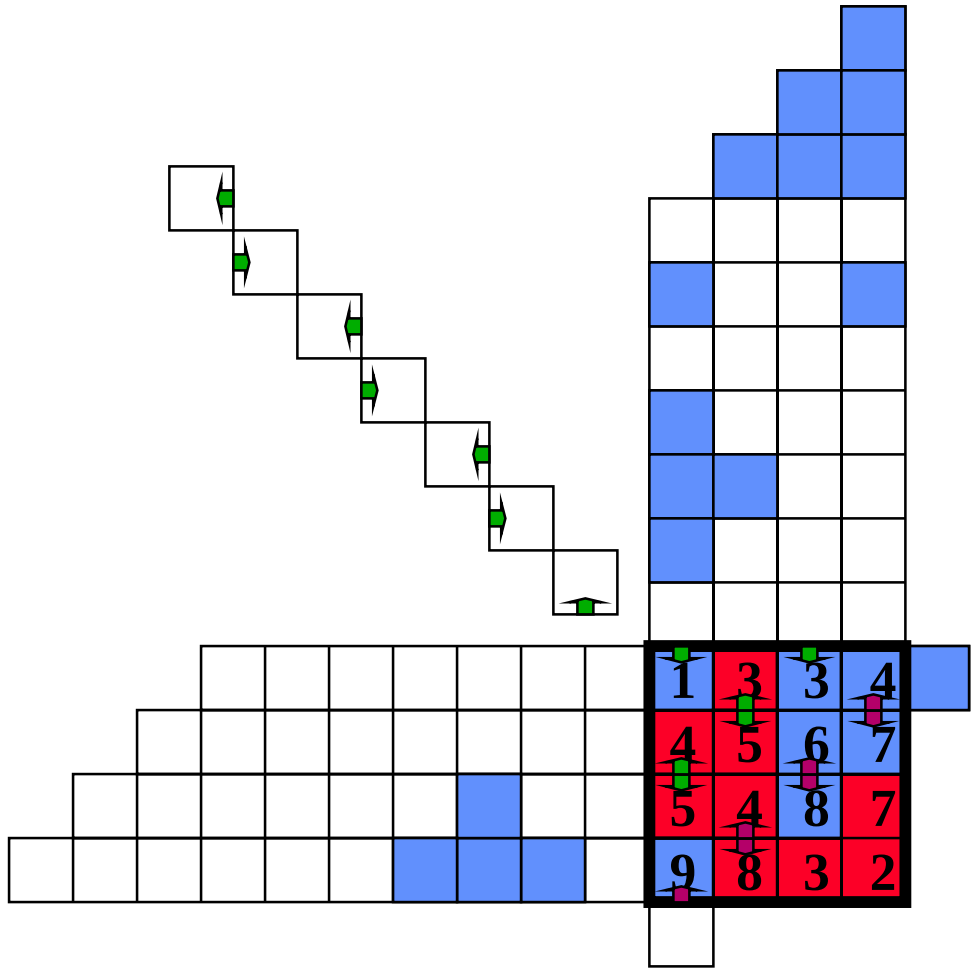

ISA merge

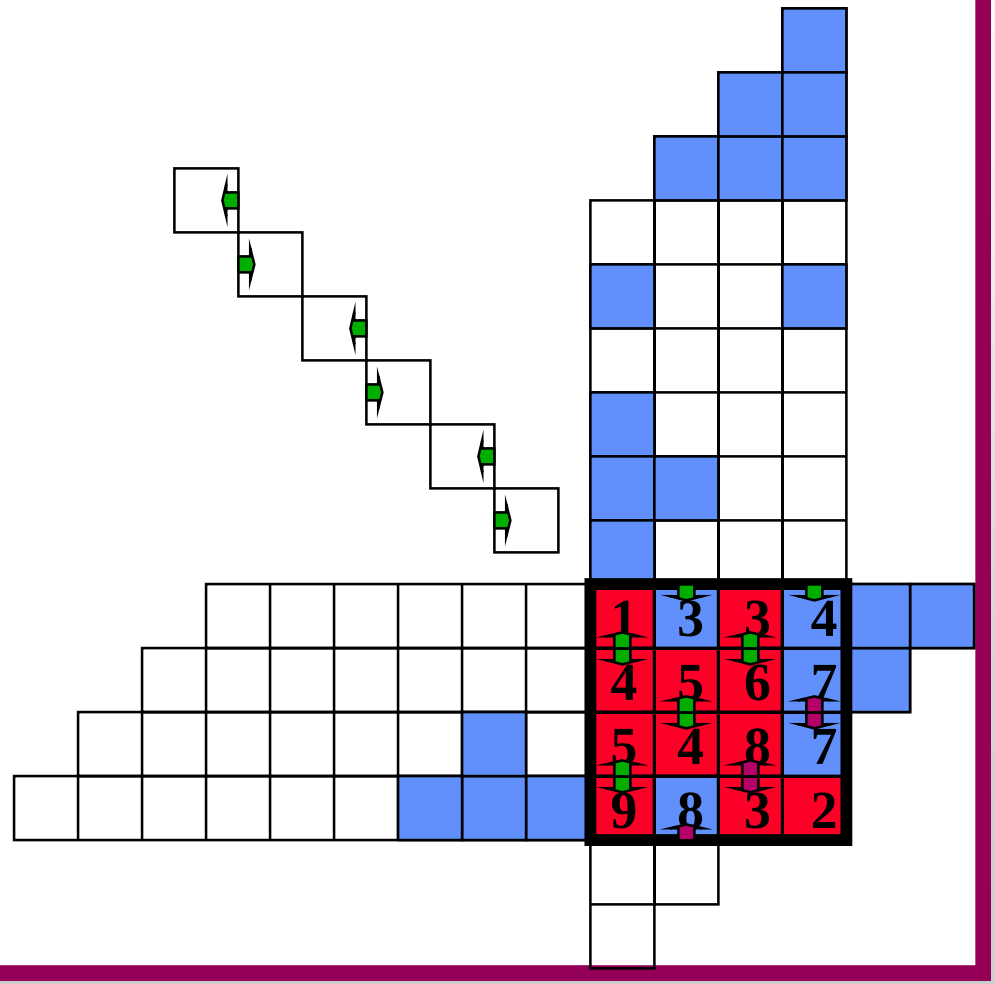

ISA merge

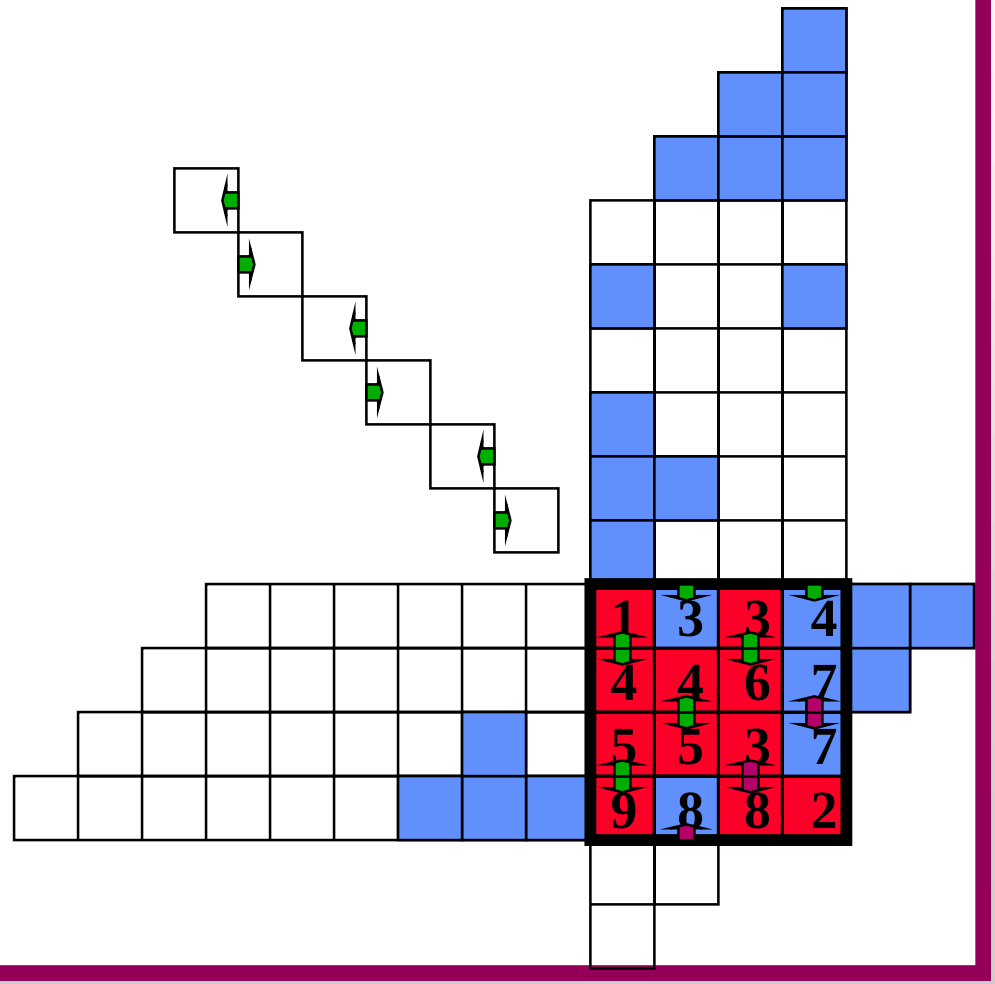

ISA merge

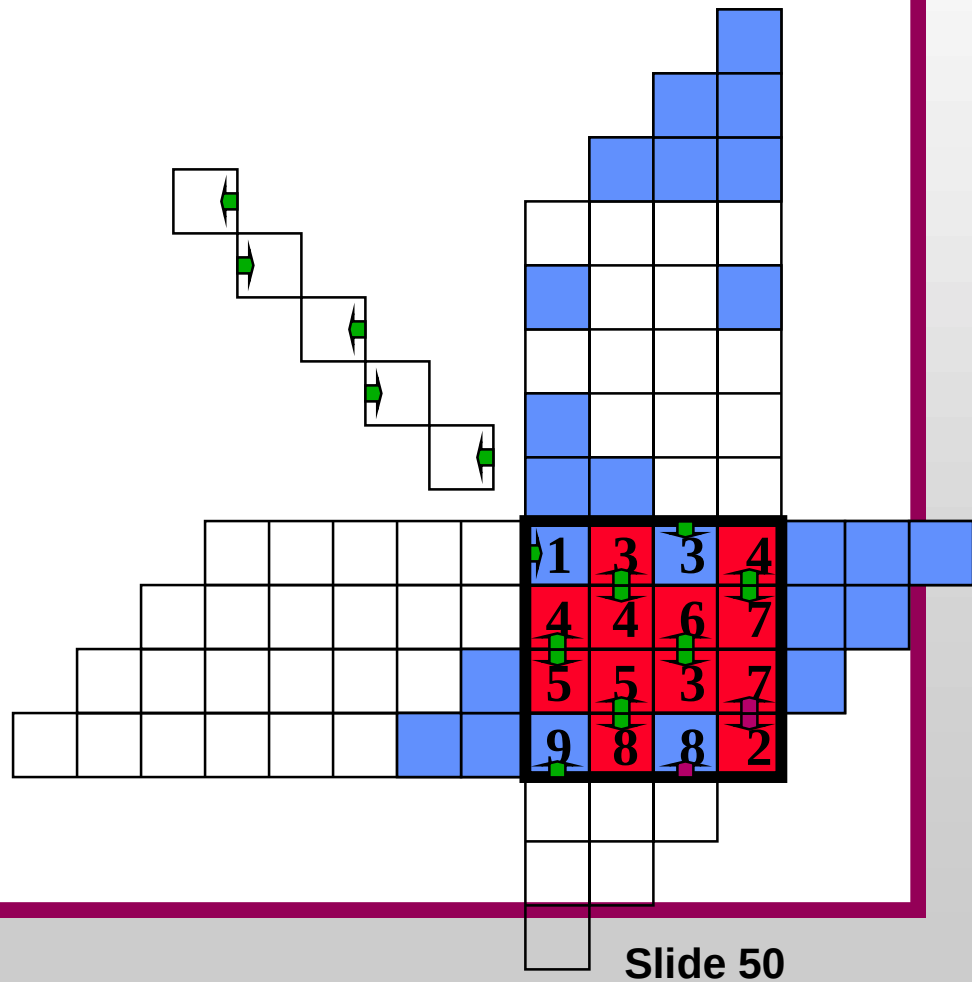

ISA merge

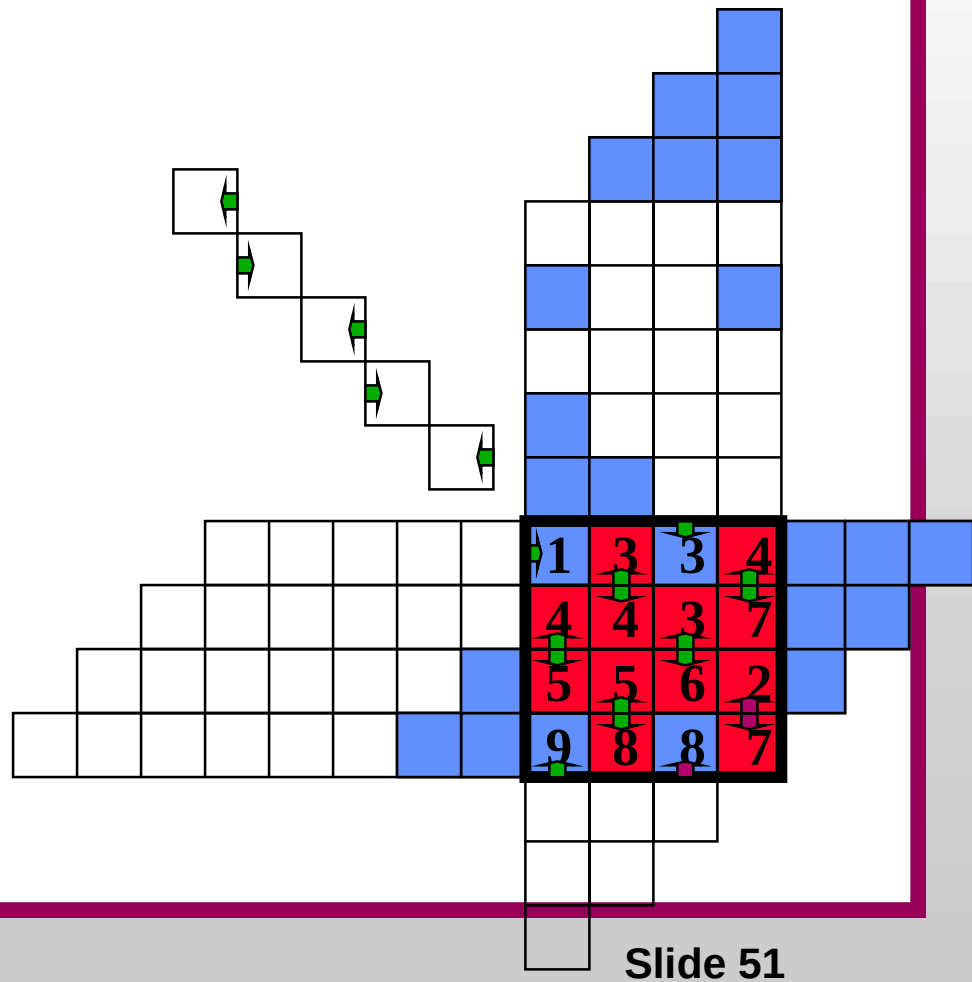

ISA merge

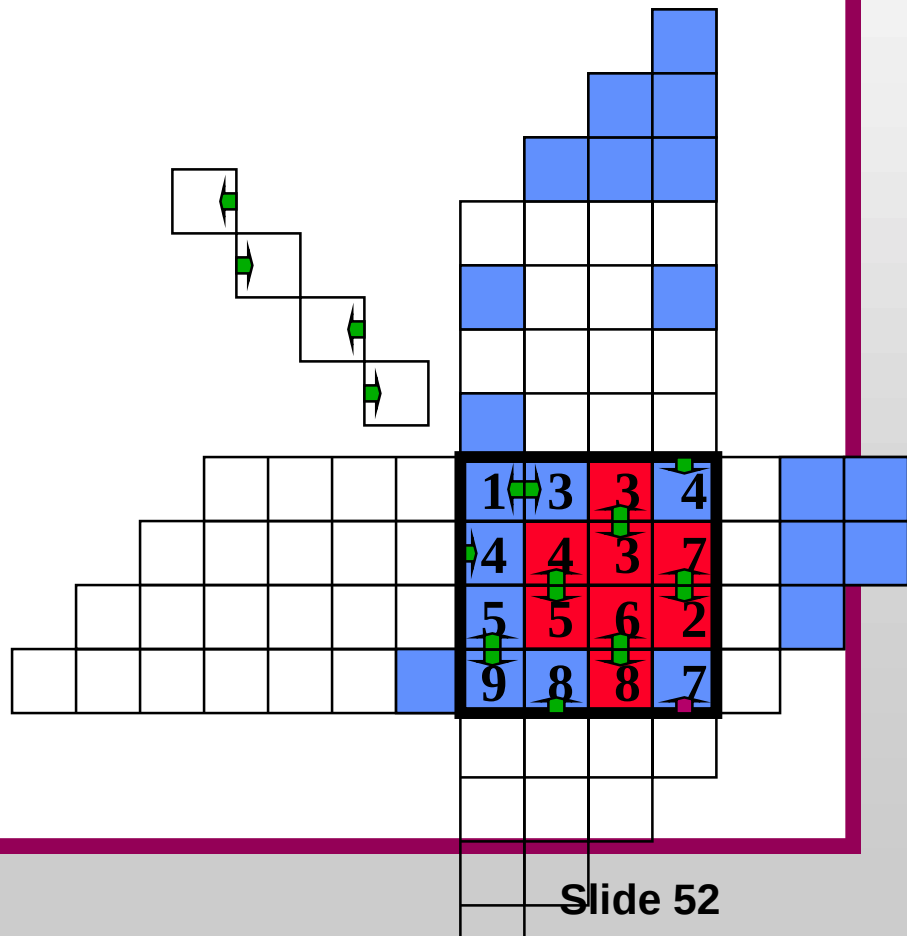

ISA merge

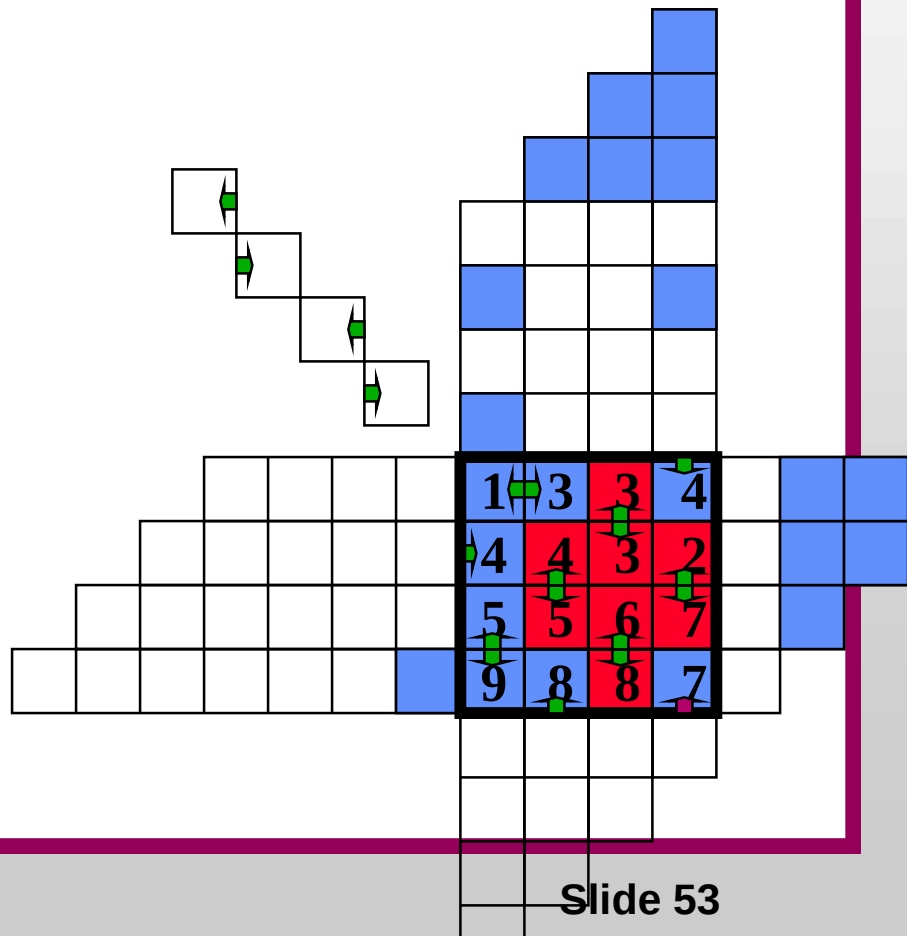

ISA merge

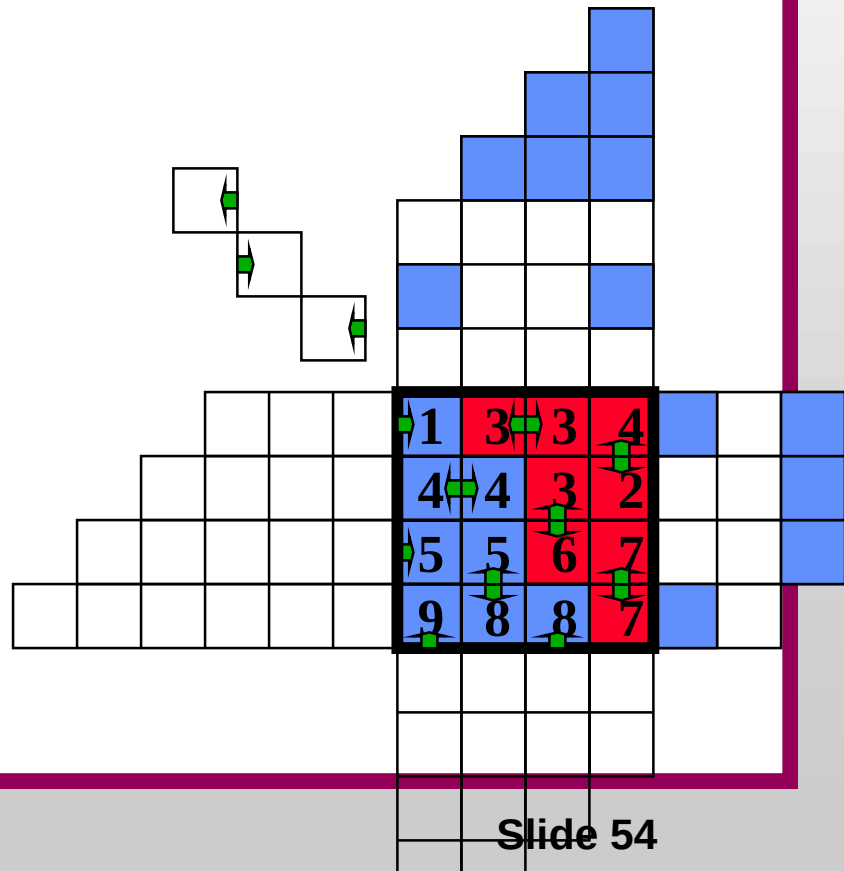

ISA merge

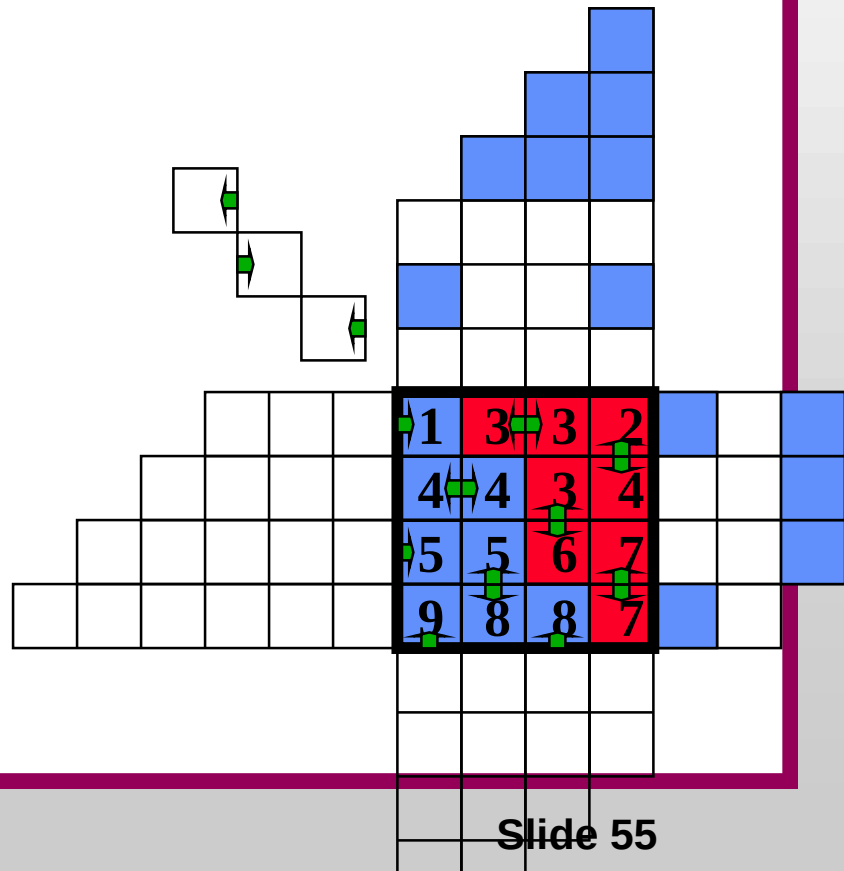

ISA merge

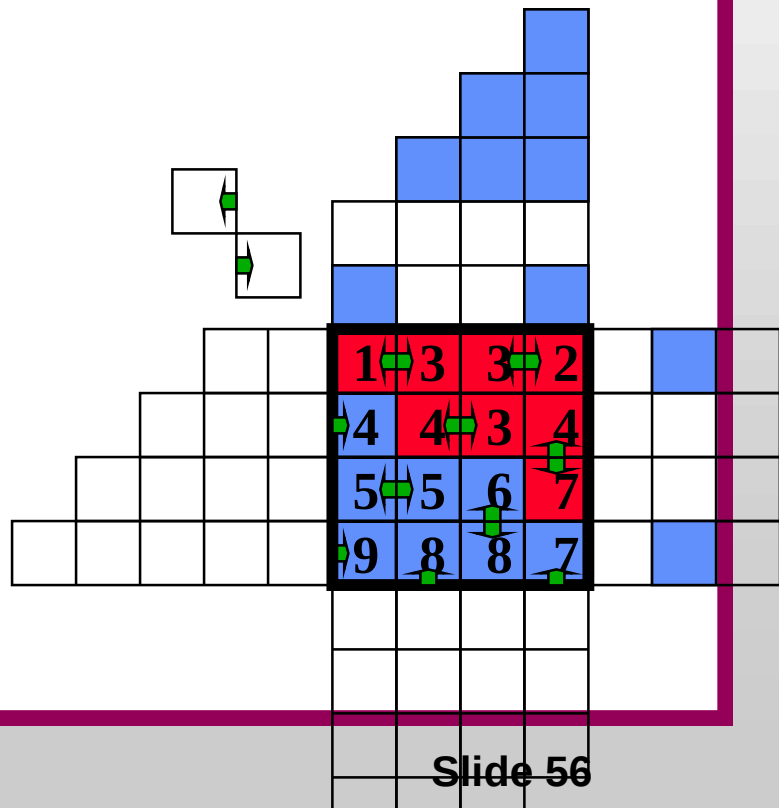

ISA merge

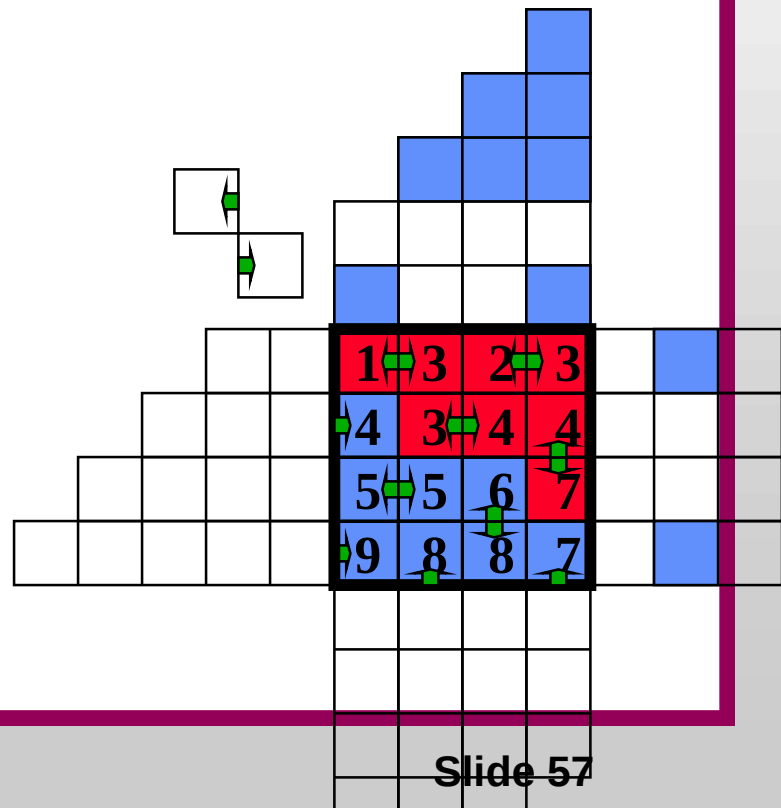

ISA merge

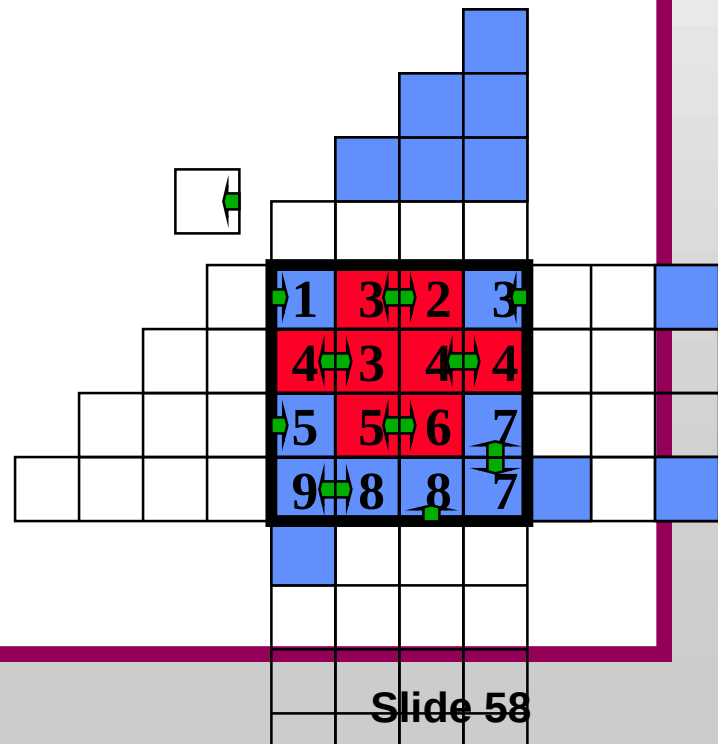

ISA merge

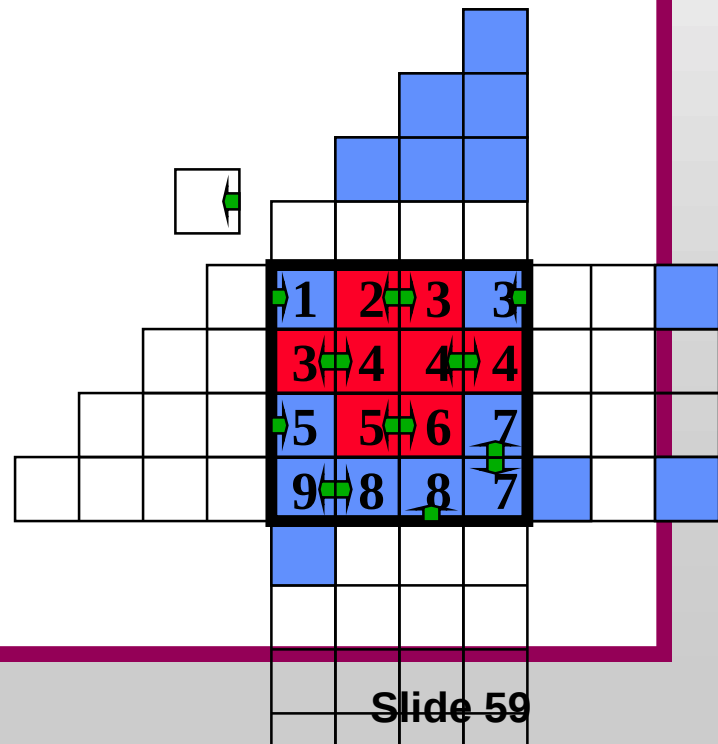

ISA merge

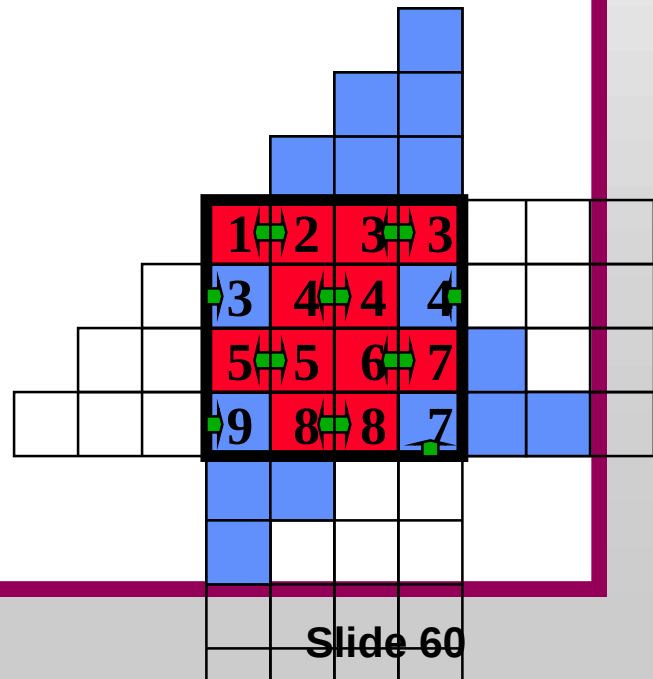

ISA merge

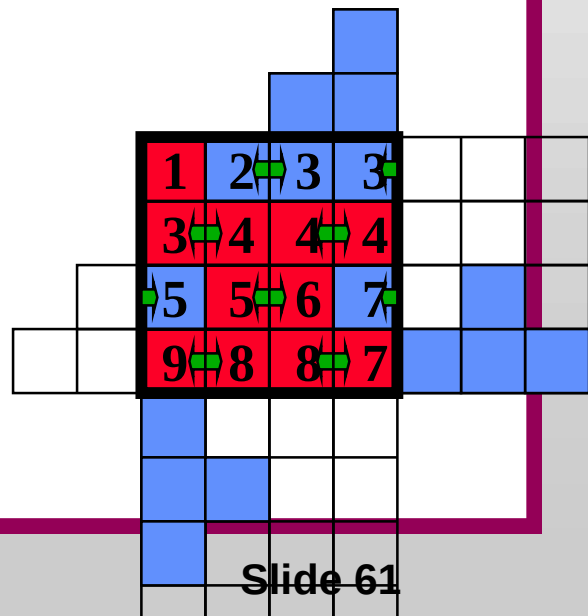

ISA merge

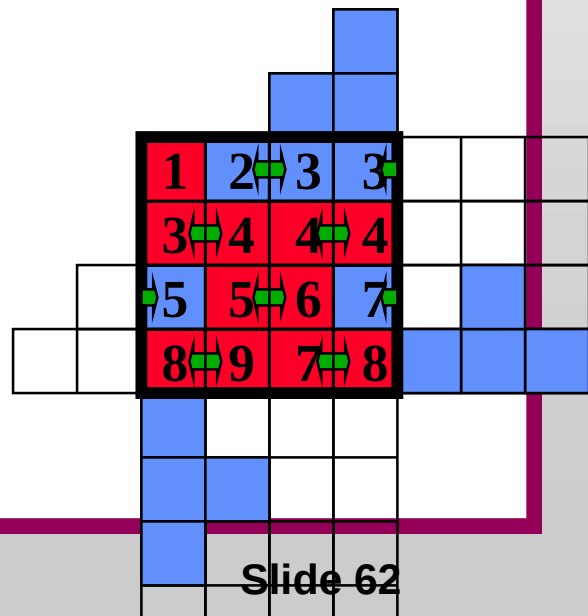

ISA merge

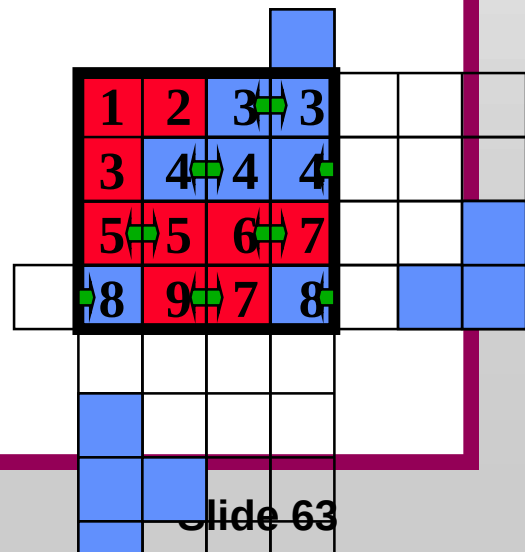

ISA merge

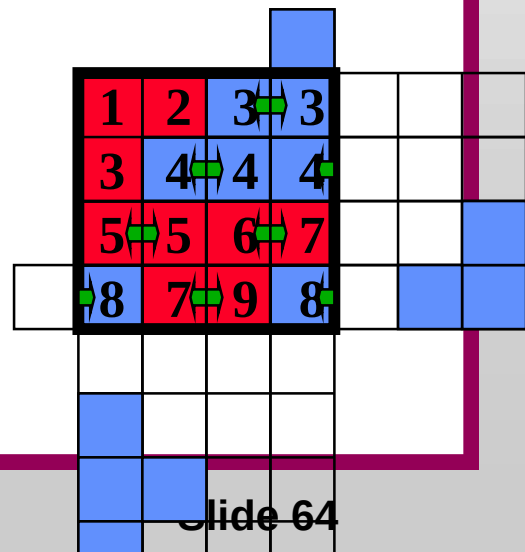

ISA merge

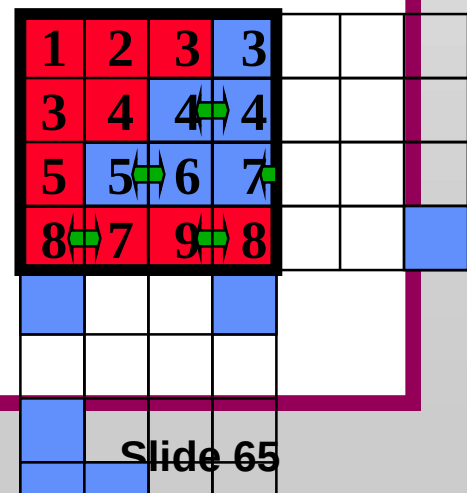

ISA merge

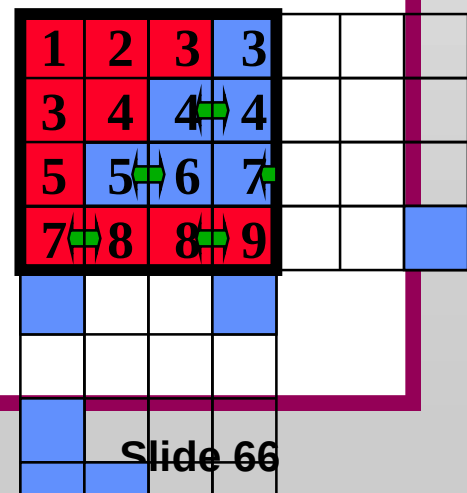

ISA merge

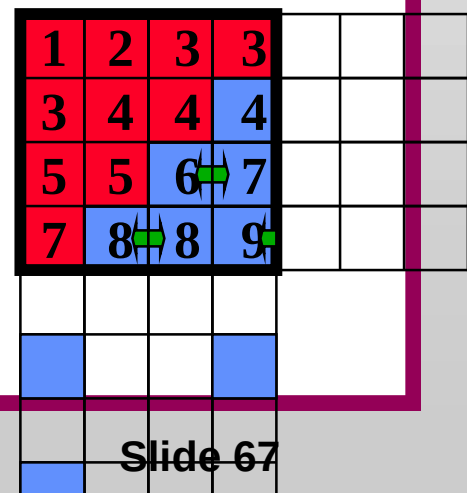

ISA merge

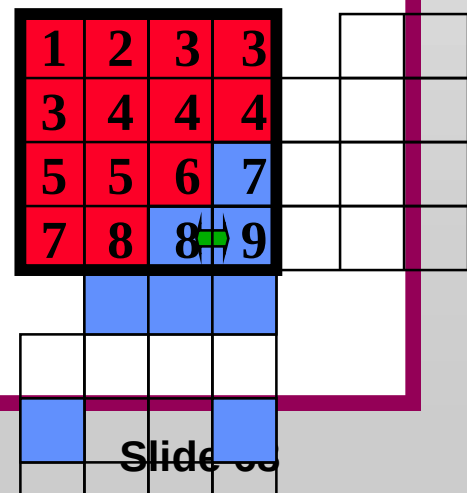

ISA merge

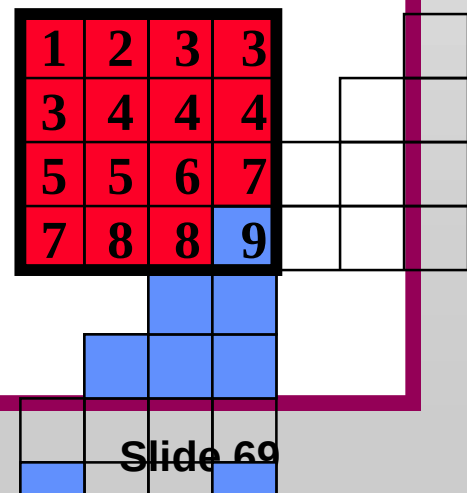

ISA merge

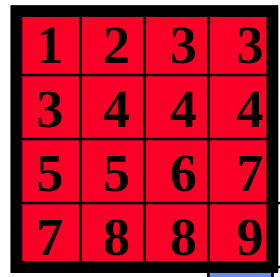

Hough transform on the ISA

- good line detection method

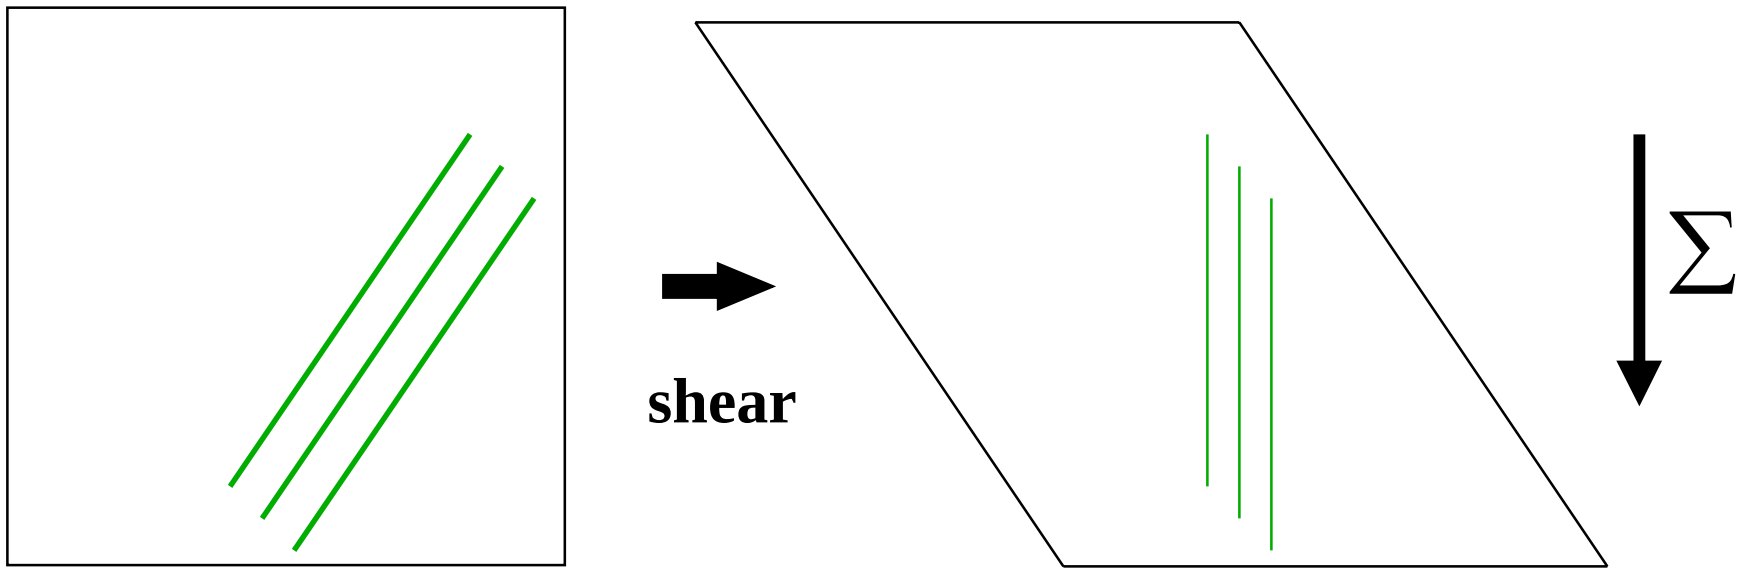

Fast tomography

robot vision

- stereo vision

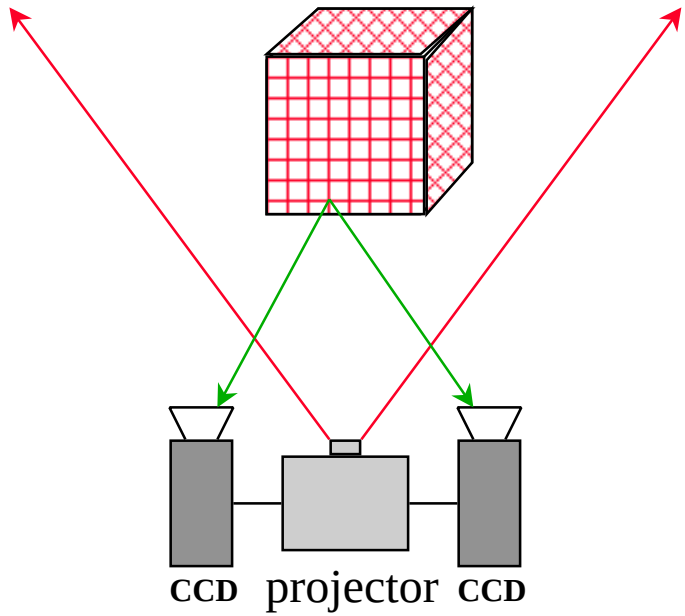

Use of the ISA

Special features:

fast aggregate functions (sum, carry)
fast local communication
no local memory
typical improvement over PC: Factor 20-30

Areas of application for ISA:

automatic optical quality control

real time signal processing
computer graphics /visualization
linear equations
Cryptography --> Tele-medicine ?

Instruction Systolic Array

1. Schröder, H., “The Instruction Systolic Array --- A Tradeoff between Flexibility and Speed”, *Computer Systems Science and Engineering*, Vol 3 No 2, April 1988
2. Schröder, H., “Top-Down Designs of Instruction Systolic Arrays for Polynomial Interpolation and Evaluation”, in *Journal of Parallel and Distributed Computing*, Vol. 6, pp 692-703, 1989
3. Kunde, M., Lang, H.W., Schimmler, M., Schmeck, H., Schröder, H., “The Instruction Systolic Array and its Relation to other Models of Parallel Computers”, *Parallel Computing*, Vol 7, pp 25-39, 1988
4. Schröder, H., Krishnamurthy, E.V., “Instruction Systolic Array Computation of the Characteristic Polynomial of a Hessenberg Matrix”, *Parallel Computing*, 17, pp 273-278, 1991
5. H. Schröder, P. Strazdins, “Programm compression on the ISA”, *Parallel Computing* 17, 207-219, 1991
6. B. Pham, H. Schröder, “An Instruction Systolic Device for Quadratic Surface Generation”, CompEuro 89, Hamburg, West Germany, May 1989
7. P. Lenders, H. Schröder, P. Strazdins, “Microprogramming Instruction Systolic Arrays”, *MICRO 22*, Dublin, August 1989
8. B. Schmidt, M. Schimmler, H. Schröder, “Morphological Hough Transform on the Instruction Systolic Array”, Euro-Par ’97 – Parallel Computing, LNCS 1300, Springer Verlag, pp. 798-806, 1997.

PIPS (1990-94)

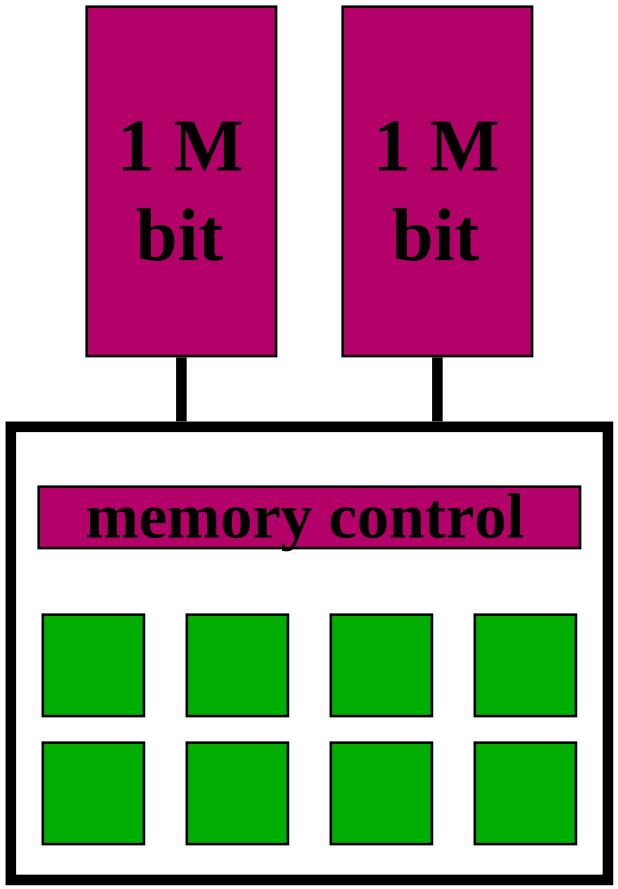

32x32 torus
16 bit parallel communication
16 bit add
prefetch

BHP -- CSIRO -- NU -- ADFA 1.4 M

PÅRC

Special features:

local memory
SIMD-torus
memory pre-fetch

Applications:

visualization
3D-simulation (CFD, FEM)

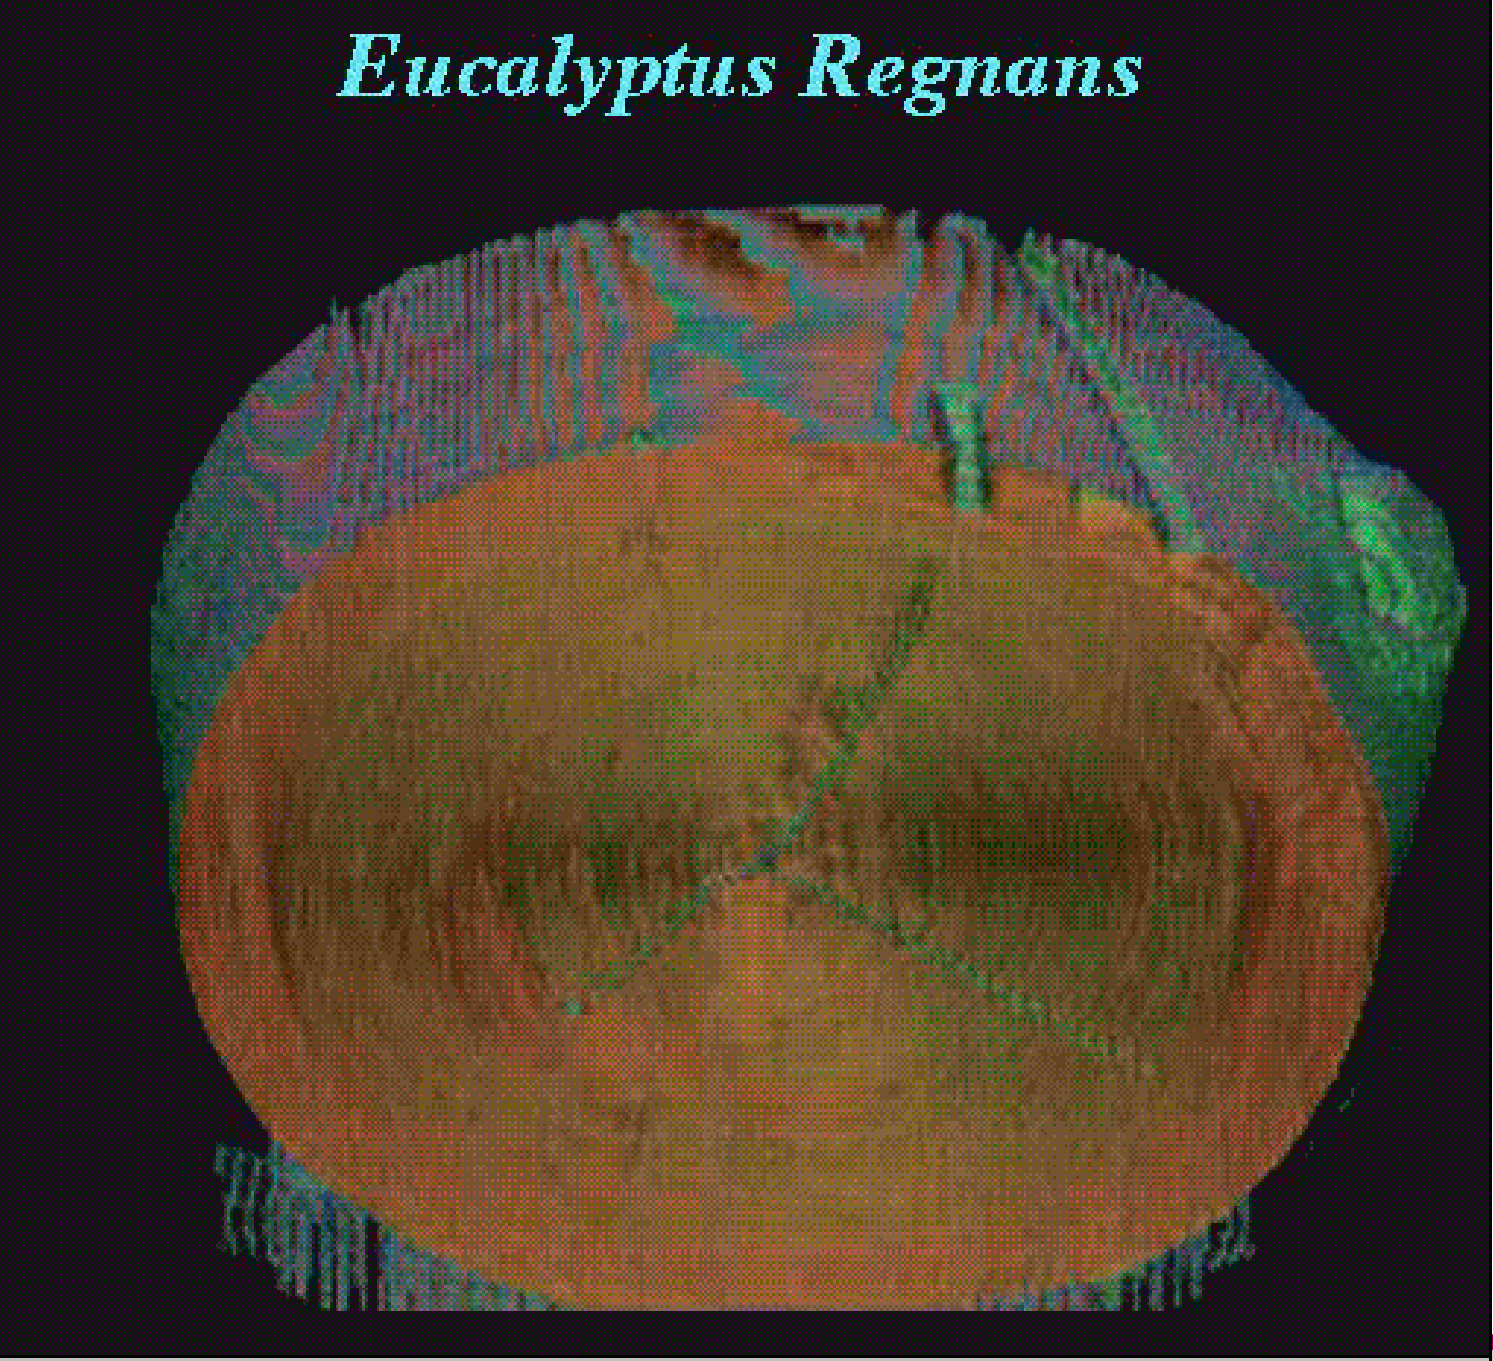

Regnans.mpg

PIPS

1. A. Spray, K.T. Lie, H. Schröder, “Test Strategies Employed in a Massively Parallel Visualization Engine”, PRFTS, Melbourne, December 1993
2. A. Spray, H. Schröder, K.T. Lie, “A Low-Cost Machine for Real-Time Visualization”, 5th International Symposium on IC Technology, Singapore, 1993
3. H. Schröder et al, “PIPADS: A Vertically Integrated Parallel Image Processing and Display System”, 5th Australian Supercomputing Conference, Melbourne, 1992.
4. R.Lang, E.Plesner, H.Schröder, A.Spray, “An efficient systolic architecture for the one-dimensional wavelet transform”, in Wavelet Applications, Harold H. Szu, Editor, Proc. SPIE 2242, p925-935, 1994.
5. R. Lang, A. Spray, H. Schröder, “2D wavelet transform on a SIMD torus of scanline processors”, *Australian Computer Science Communications*, Vol 17, Nr 17, 271-277, 1995.
6. P. Bray, S.W. Chan, K.T.Lie, Meiyun, H. Schröder, A. Spray, “A torus for scan-line based image processing”, Presented at the Post-ISCA Special Purpose Architectures Workshop, May 1992.
7. A. Spray, H. Schröder, K.T. Lie, E. Plesner, P. Bray, “PIPADS --- A Low-Cost Real-Time Visualization Tool”, Supercomputing, Melbourne, 1993.
8. H. Schröder, A. Spray, “PIPS a massively parallel image processing system”, Workshop PARAGRAPH'94, Linz, March 1994.
9. H. Schröder, “3D-Visualisation based on Scan-line Image Processing”, invited speaker, Workshop MWTAI’97 in Missen-Wilhams, March 1997

Use in industry ?

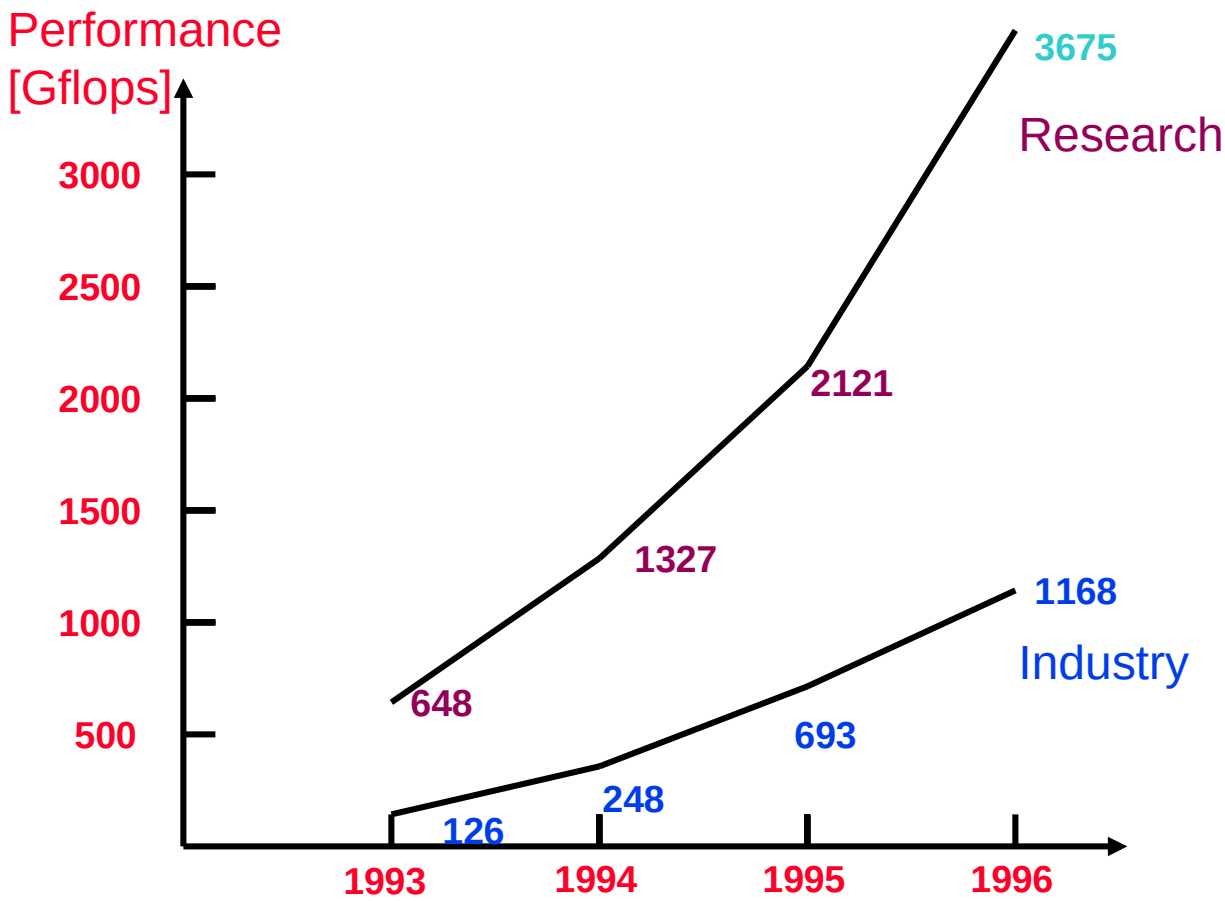

Investments

Investments into parallel computers
[M$]

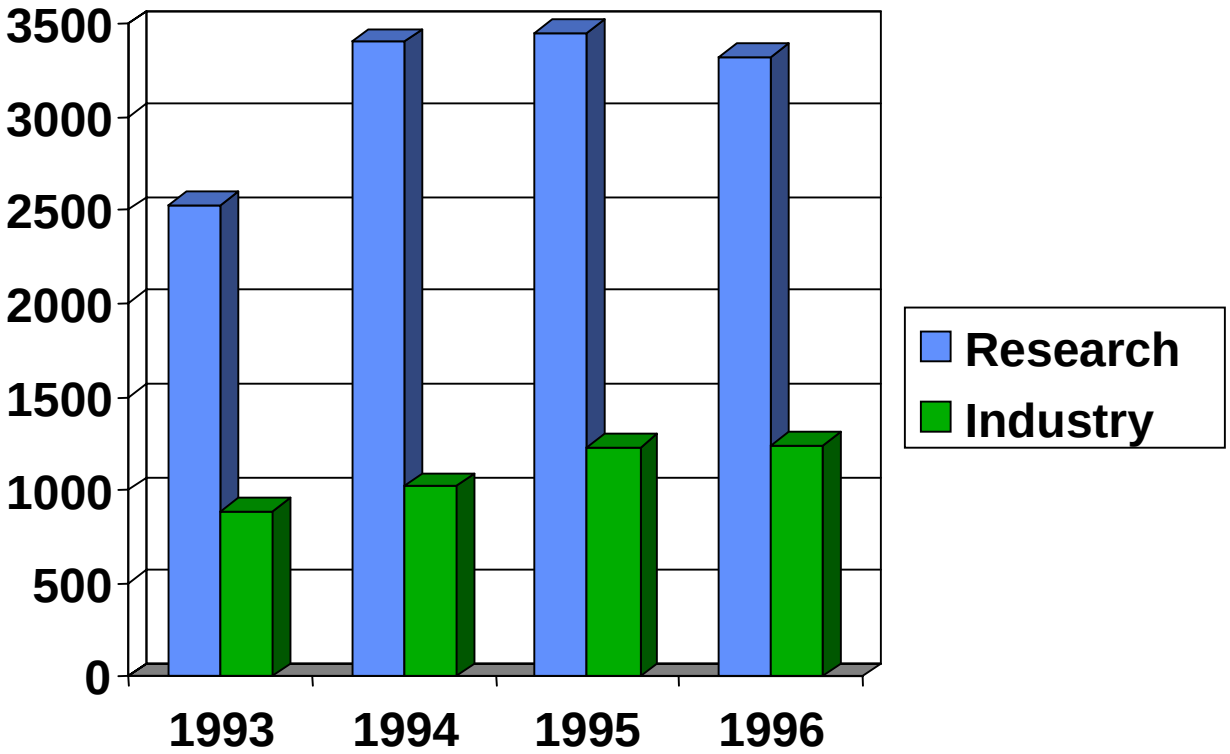

Concentration

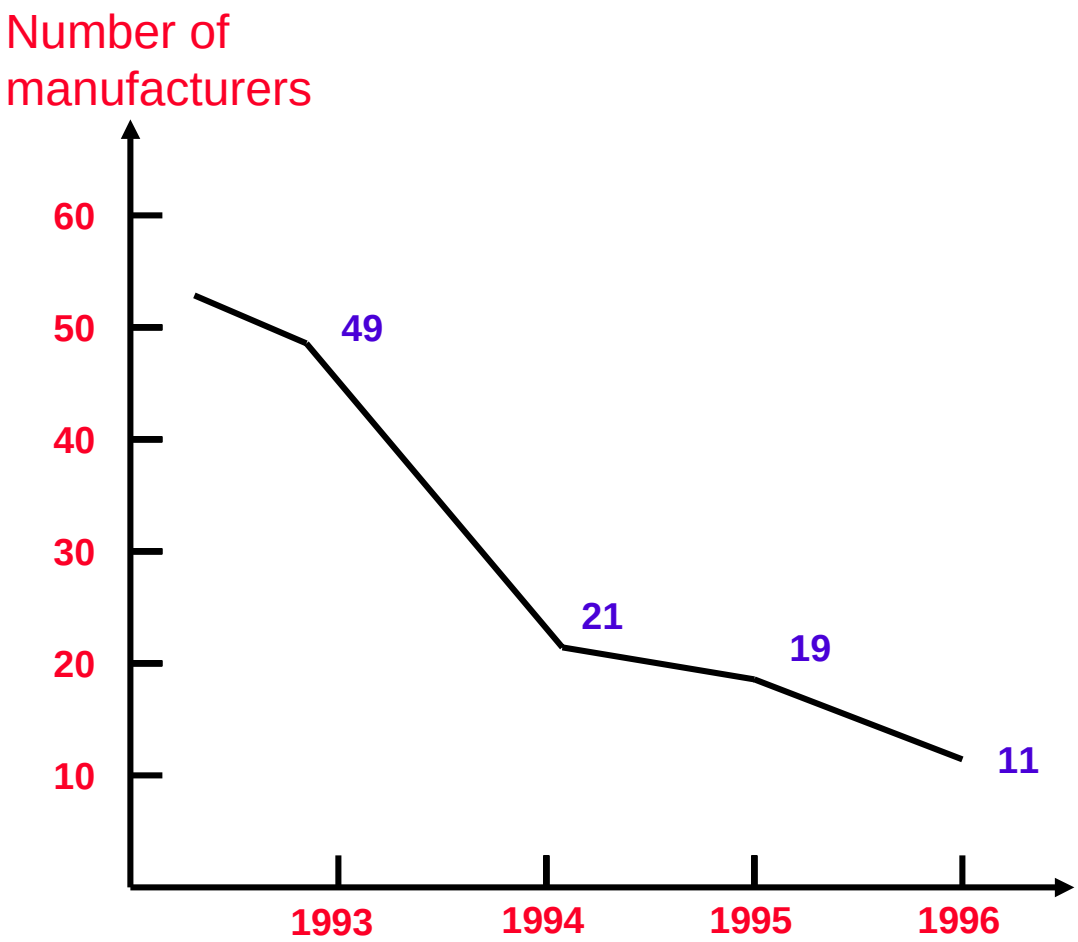

Degree of Parallelism

Number of new Systems

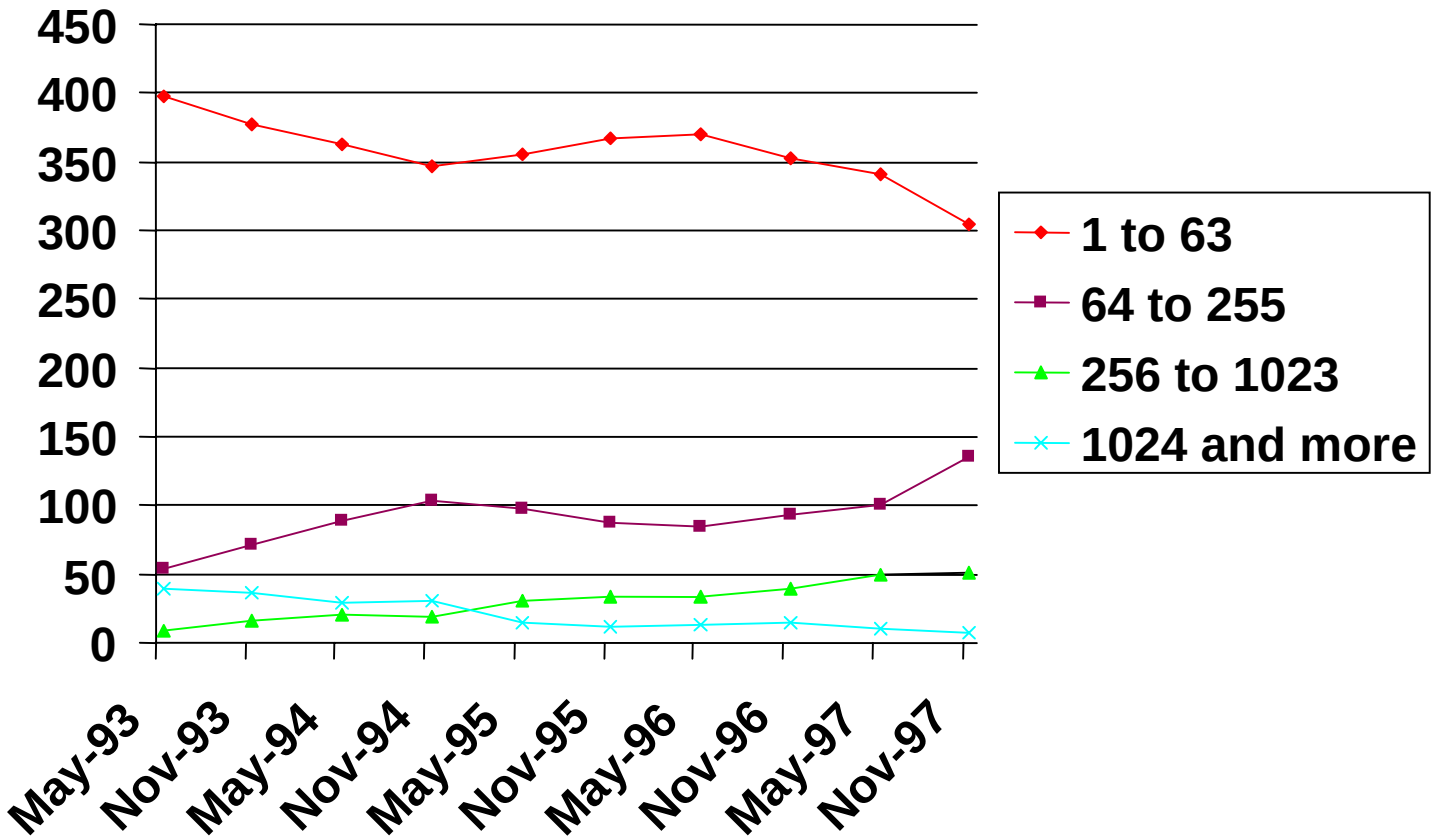

Evaluation

Cost * computation time

- Parallel computers with standard components
- Imbedded parallel systems

reconfigurable mesh

reconfigurable mesh =
mesh + interior connections

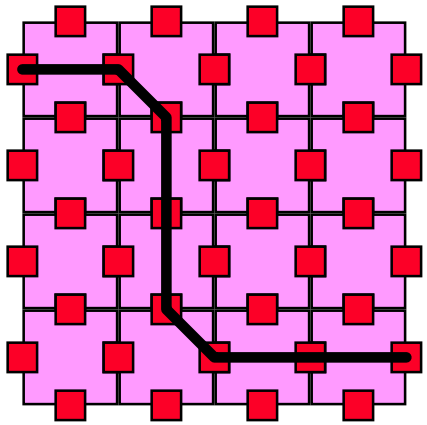

low cost

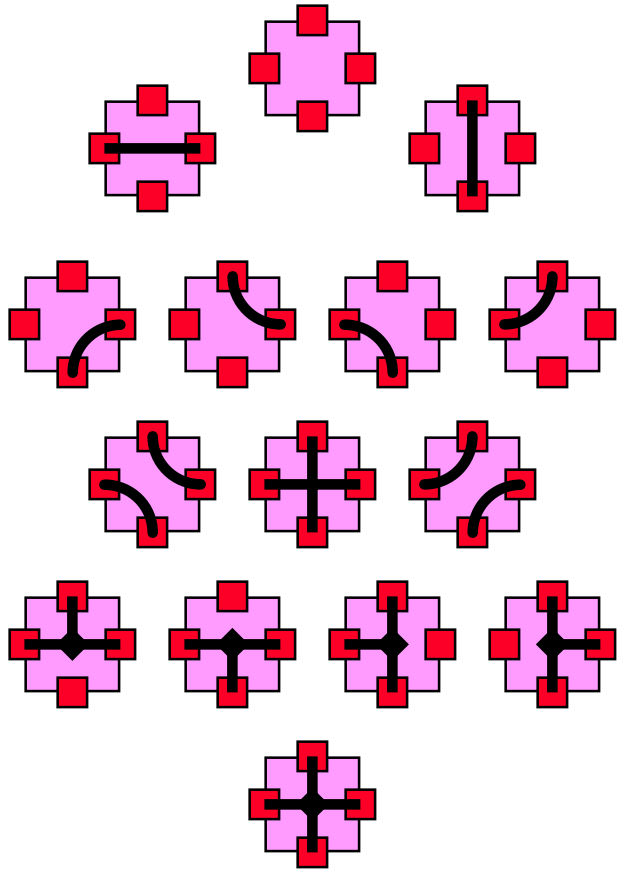

15 positions

global OR and modulo 3

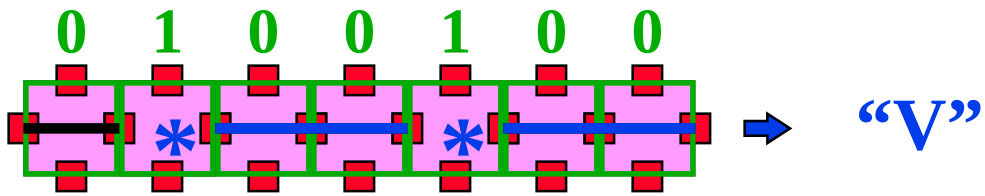

log n on EREW-PRAM

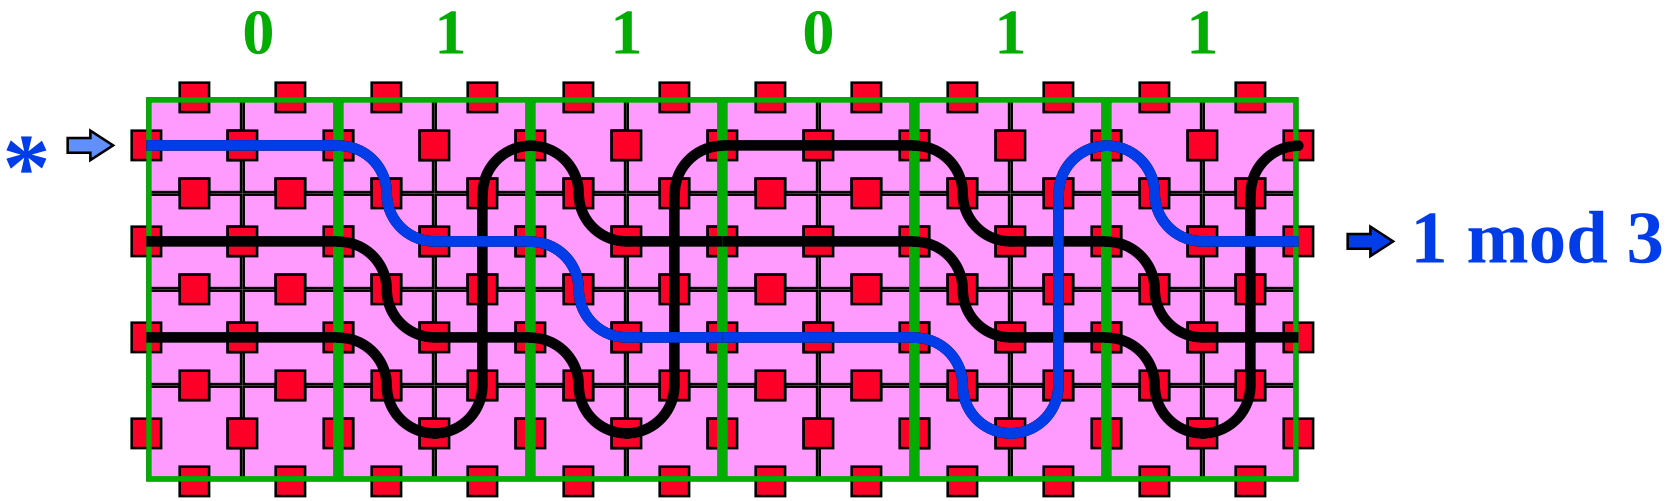

log n / log log n on CRCW-PRAM

sorting with all-to-all mapping

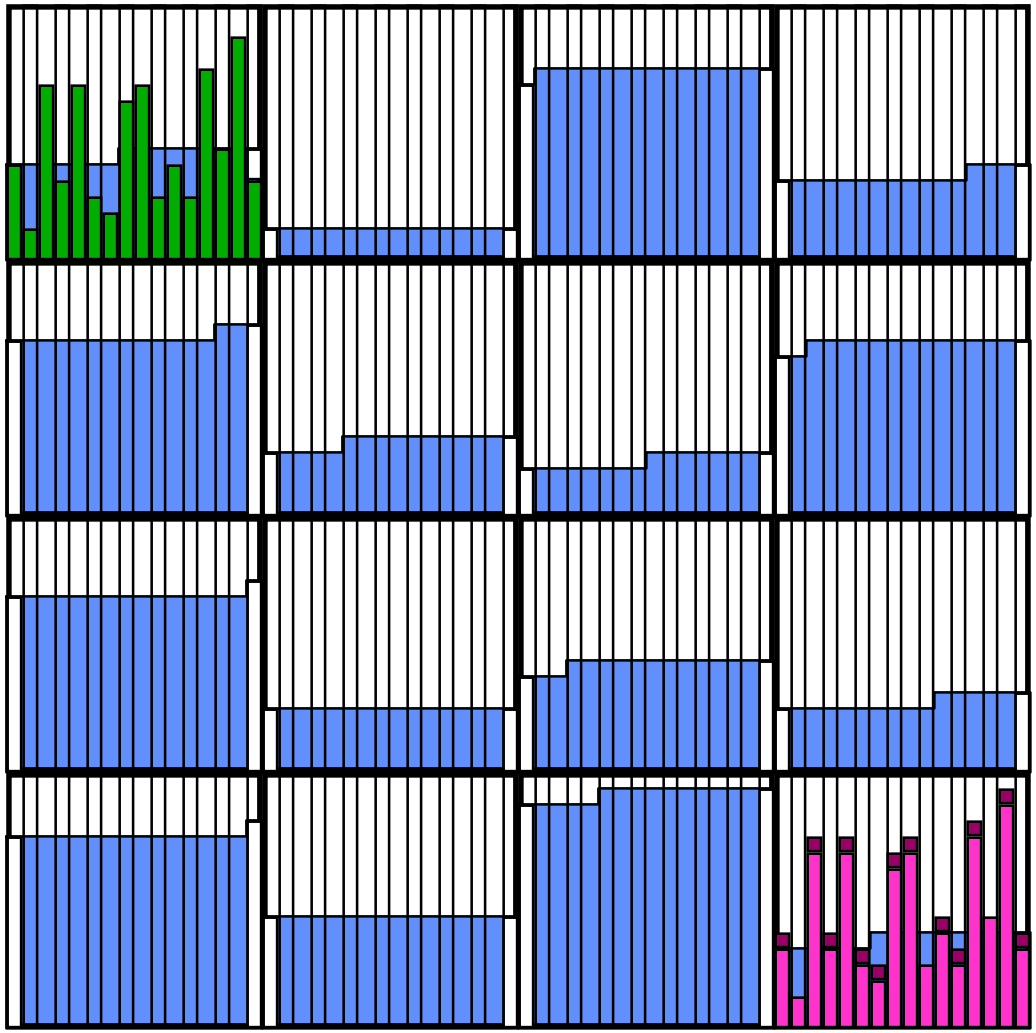

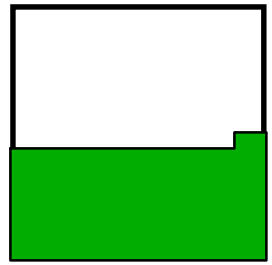

Sorting:

sort blocks

all-to-all (columns)

sort blocks

all-to-all (rows)

o-e-sort blocks

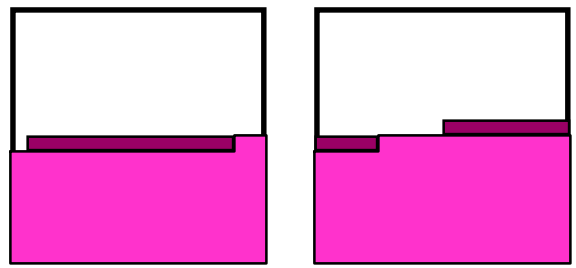

all-to-all mapping

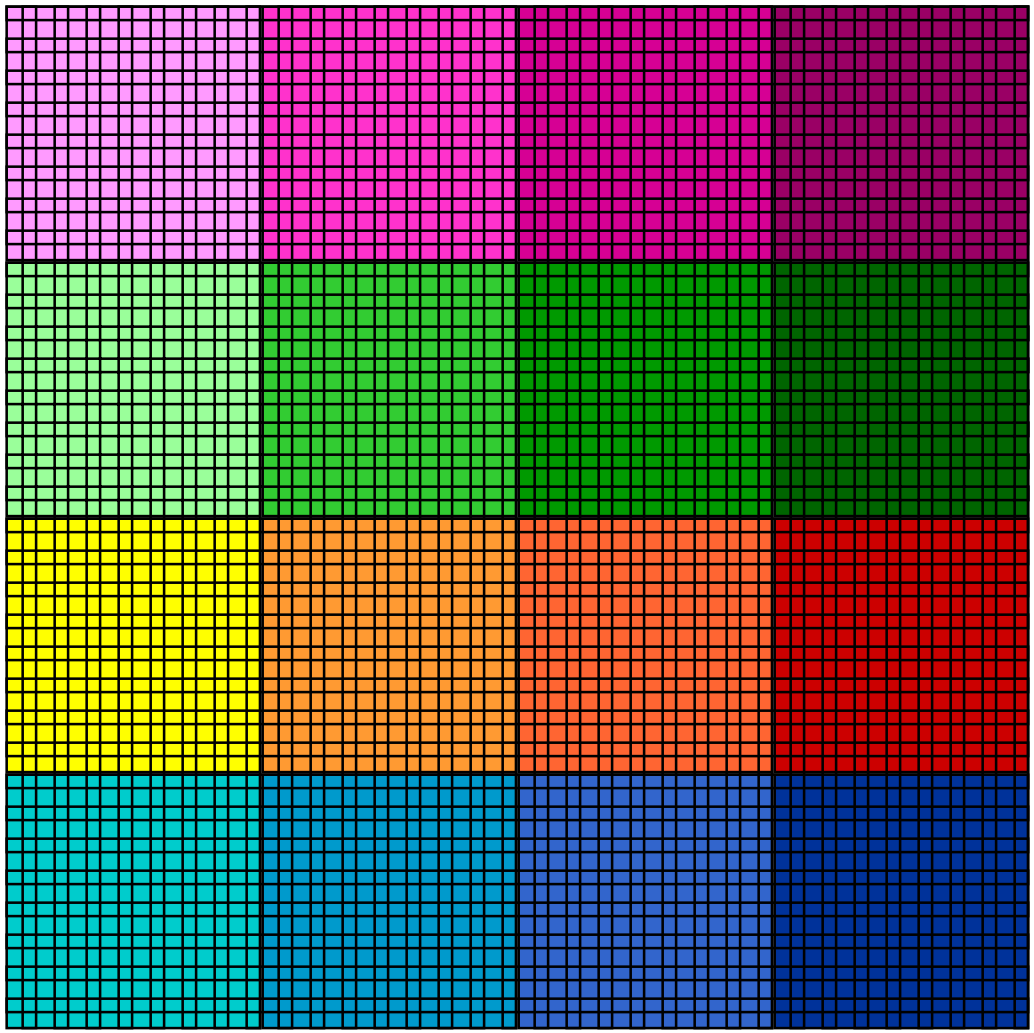

$n^{2/3}$

n x n

vertical all-to-all

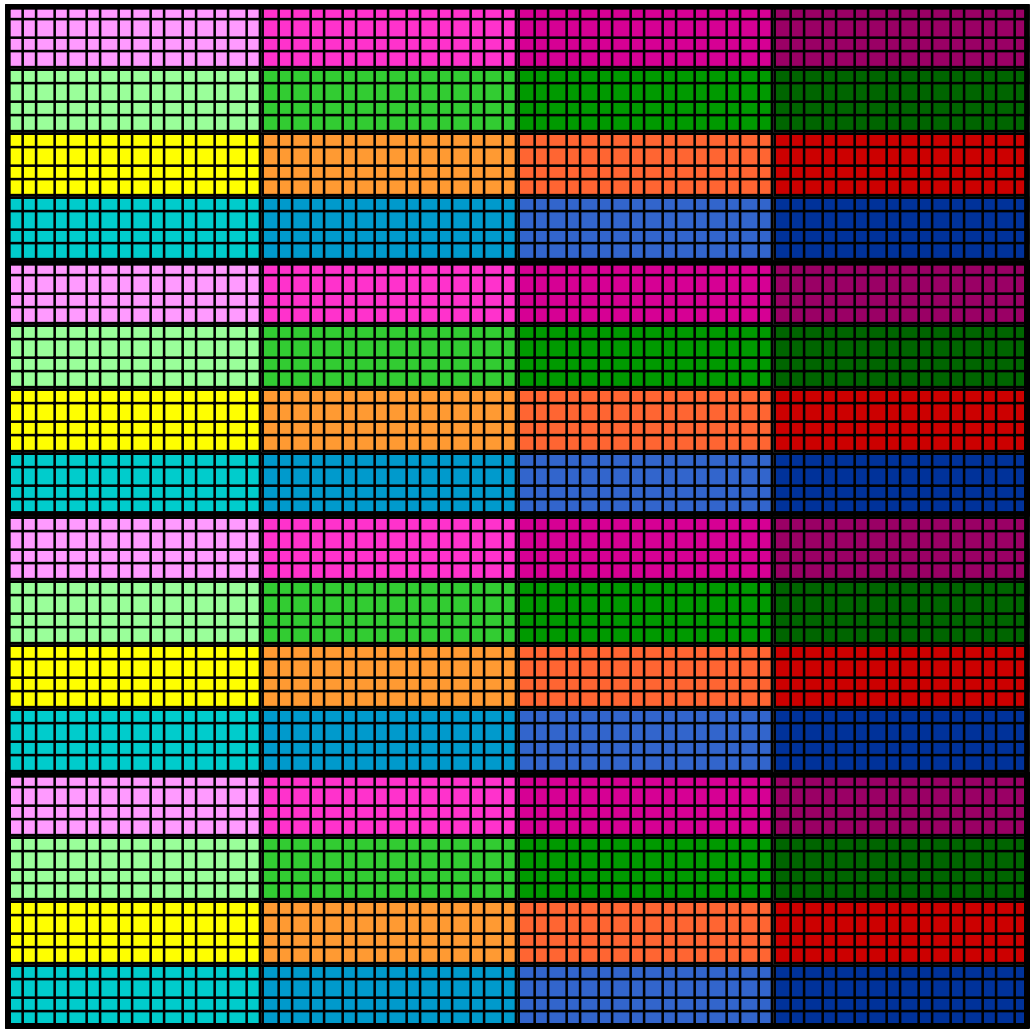

horizontal all-to-all

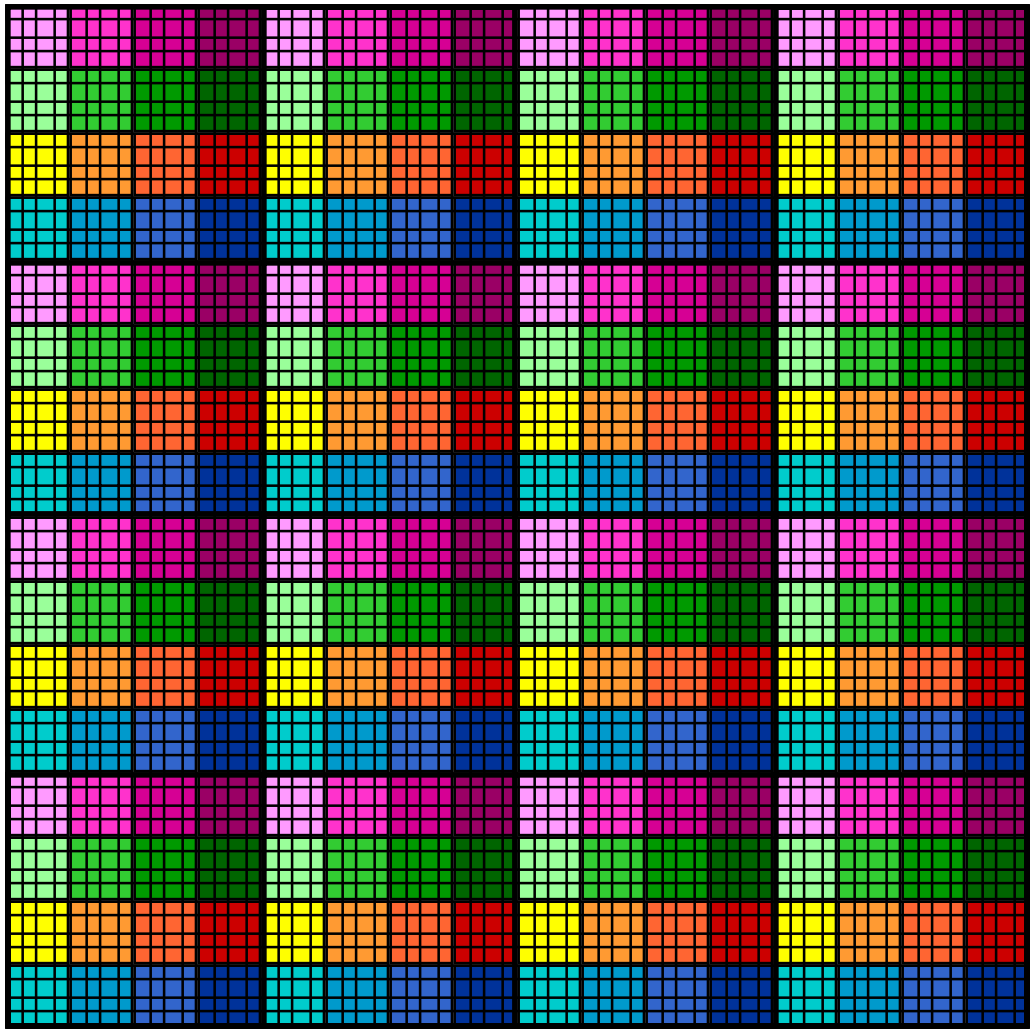

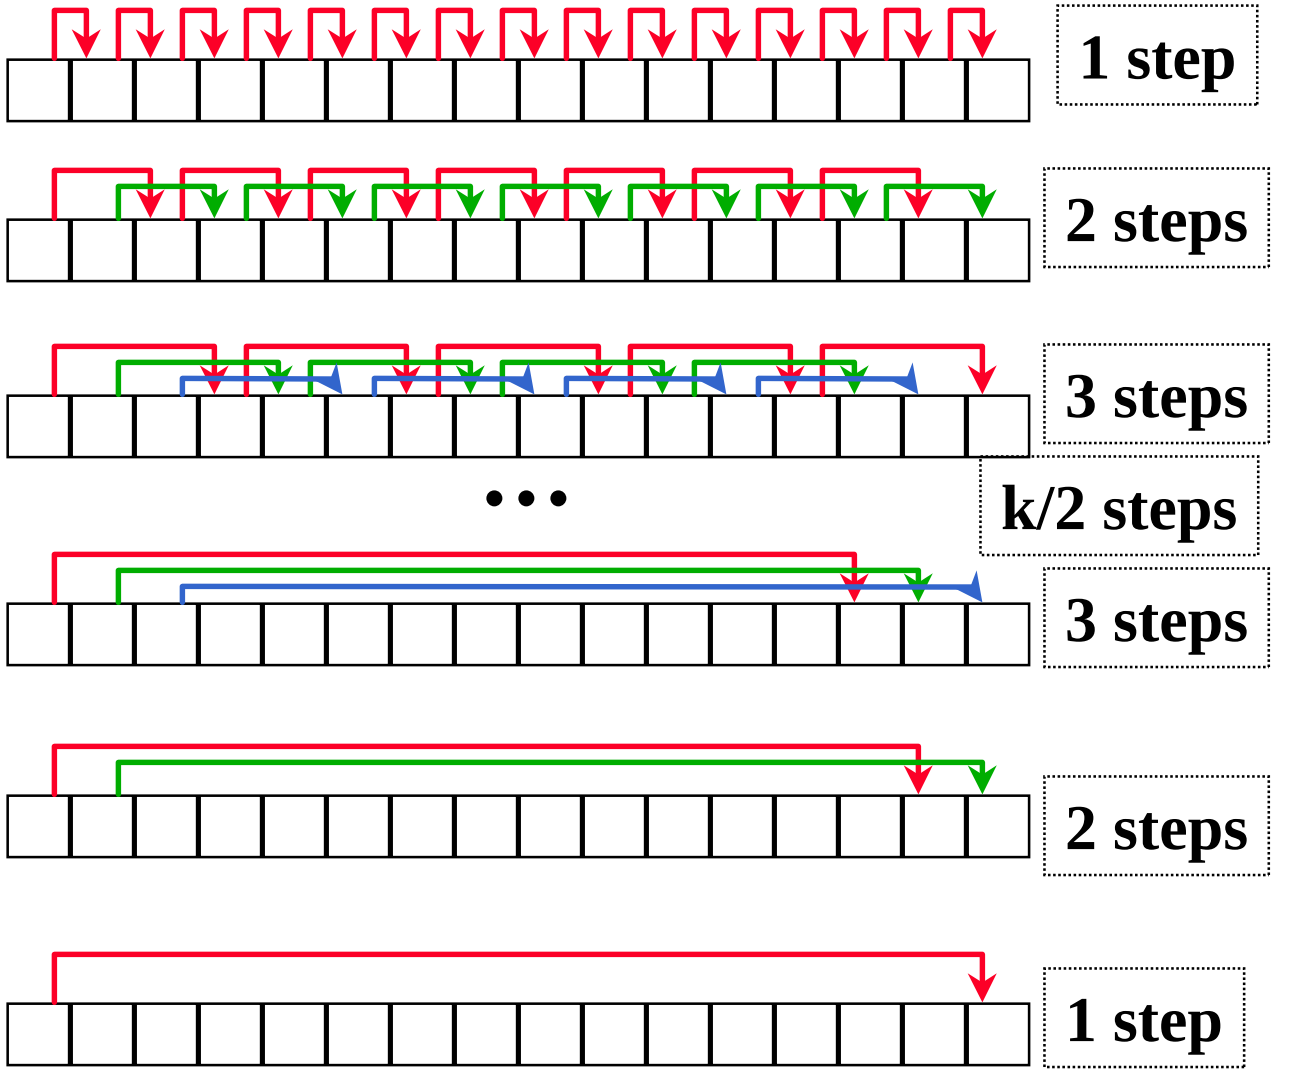
1 step
(k/2)² steps
2 steps
3 steps
k/2 steps
3 steps
2 steps
1 step

sorting in optimal time

$(k/2)^2$ steps
$k=n^{1/3}$
each step takes $n^{1/3}$ time
--> T= n/4


x 2

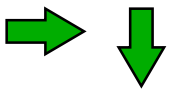
x 2

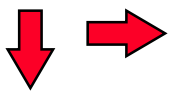
/2

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

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)

Reconfigurable mesh

Special features

SIMD

constant diameter

faster than PRAM ?

Suitable applications

routing/sorting/load balancing

sparse matrix multiplication

segmentation / component labeling

feature extraction

image database ?

Reconfigurable mesh

1. Kapoor, H. Schröder, B. Beresford-Smith, “Connected Component Labelling on 3D Reconfigurable Mesh Architecture”, *Parallel Processing*, 1994
2. M. Kaufmann, H. Schröder, J. Sibeyn, “Routing and Sorting on Reconfigurable Meshes”, *Parallel Processing Letters*, Vol 5, Nr 1, pp 81-96, 1995
3. G. Turner, H. Schröder, “Token Distribution and Load Balancing on reconfigurable, d-D. Meshes”, *Journal of Parallel Architectures and Algorithms*, 1996
4. M. Middendorf, H. Schmeck, H. Schröder, G. Turner. “Multiplikation of matrices with different sparseness properties on dynamically reconfigurable meshes”, to appear in VLSI Design.
5. Kapoor A., Schröder H., Beresford-Smith B., “Constant Time sorting on a Reconfigurable Mesh “, Australian Computer Science Conference, February 1993.
6. D. Yu, H. Schröder, “Parallel Border following, Connecting and Labeling for Binary Image with 8-neighborhood Coding on a Mesh with Bi-reconfigurable Buses”, 3rd International Workshop on Parallel Image Analysis, Maryland, 1994.
7. Kapoor, H. Schröder, B. Beresford-Smith, “Deterministic Permutation Routing on the Reconfigurable Mesh”, Proceedings 8th IPPS, pp 536-540, Mexico, 1994.
8. G. Turner, H. Schröder, “Fast Token Distribution on Reconfigurable Meshes”, in Proceedings of the 2nd Australasian Conference on Parallel and Real-Time Systems (PART)}, vol 2, pp 127-133, 1995
9. H. Schmeck, H. Schröder, G. Turner, “Efficient Sparse Matrix Multiplication on a Reconfigurable Mesh”, PARS-Workshop, Stuttgart, Germany, October 1995.
10. H. Schröder, “Mathematical Morphology for Robot Vision on a Reconfigurable Mesh Architecture”, ISCA Special Purpose Architectures Workshop, May 1992.
11. M. Kaufmann, H. Schröder, J. Sibeyn, “Asymptotically Optimal and Practical Routing on the Reconfigurable Mesh”, GI/ITG-Workshop “Architekturen für hochintegrierte Schaltungen”, Schloß Dagstuhl, July1994.

Optical Highway

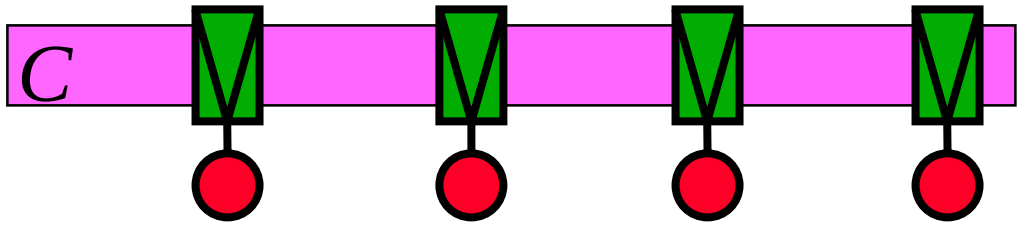

All-to-all connection

$$P^3 W = C$$

$$P = \#(processors)$$

$$W = width$$

$$C = 10^6$$

W=1; P=100
W=100; P=22

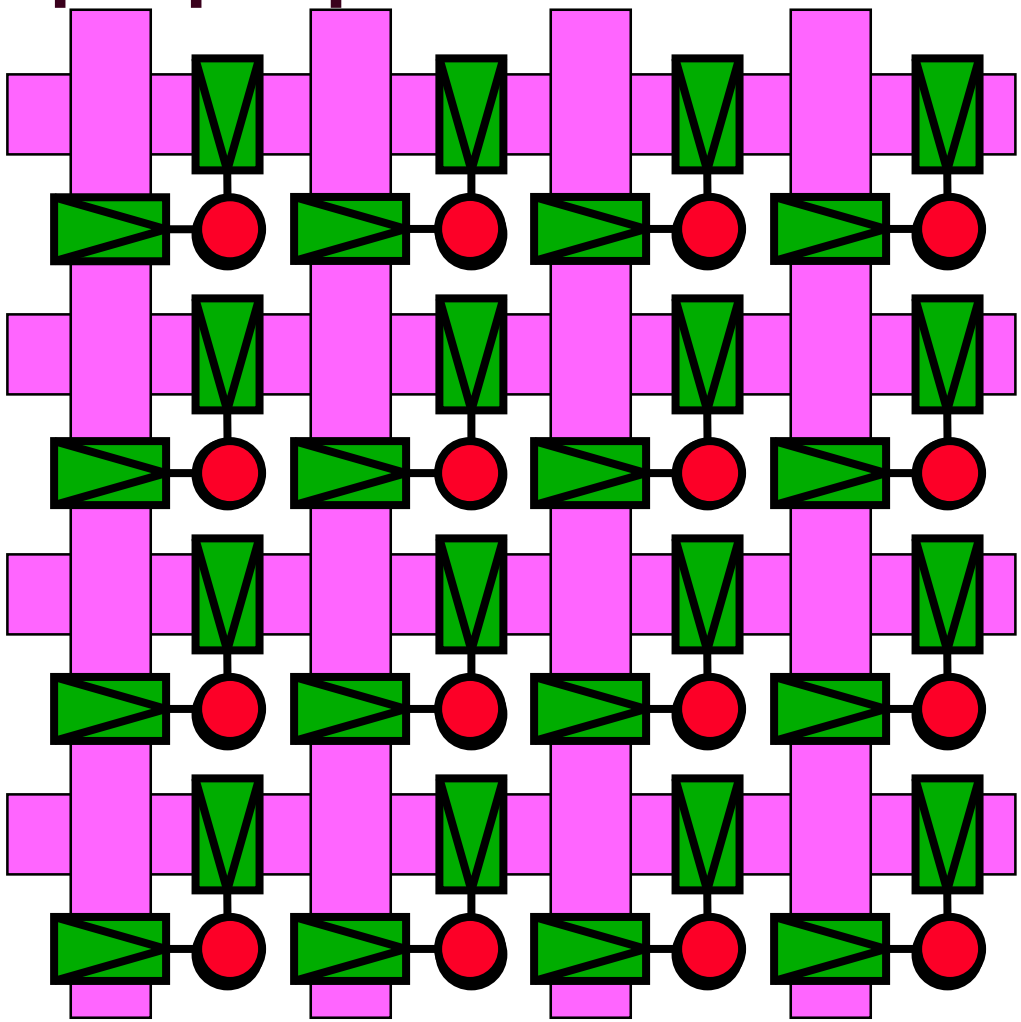

Horizontal
all-to-all

Vertical
all-to-all

Features of optically connected meshes

SIMD/SPMD/MIMD

implement all major architectures

all-to-all communication in 2 steps

Bulk synchronous processing (BSP)

no latency hiding

no pin-limitation

Applications

coarse grain parallel computing only?

ray-tracing ?

???

Optical Highway

1. H. Schröder et al, “RMB --- A Reconfigurable Multiple Bus Network”, HPCA 96, San Jose, 1996
2. H. Schröder, O. Sykora, I. Vrto, “Optical All-to-All Communication for some Product Graphs”, SOFSEM '97, Milovy, Czech Republic, 1997

Bisection-width / Diameter

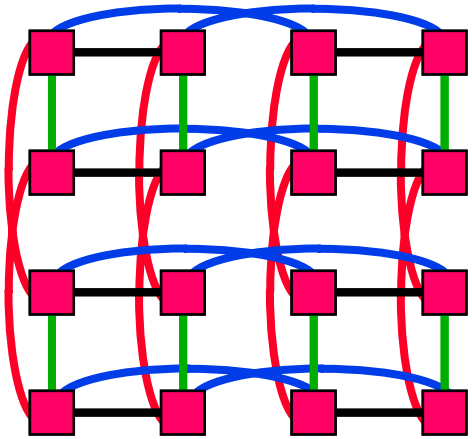

Diameter log n
bisection width n

SA
ISA
PIPS

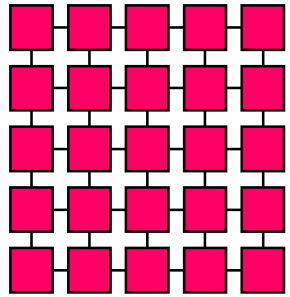

diameter $\sqrt{n}$
bisection width $\sqrt{n}$

RM

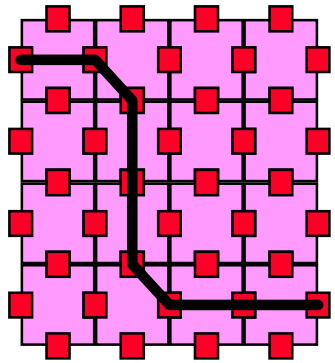

Diameter 1
bisection width $\sqrt{n}$

OH

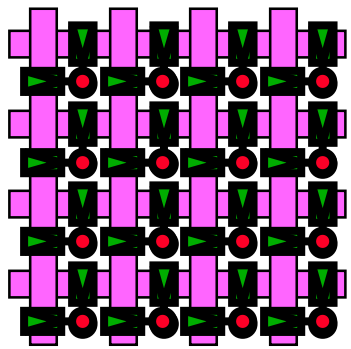

Diameter 1
bisection width n

bisection width n

P Å R C

Suitable problems ?

SA
ISA
PIPS

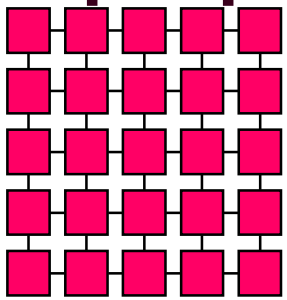

SA: suitable applications?

ISA: 2D-problems, aggregate functions
local communication

PIPS: 3D-problems, local communication

RM

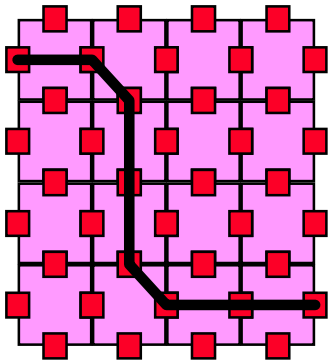

RM: diameter-bound > bisection-width-bound

OH

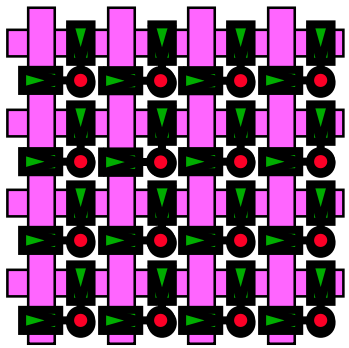

OH: PRAM equivalent?


STOP