
Systolic Arrays and Algorithms


Systolic Arrays

- This is a form of **pipelining**, sometimes in more than one dimension.
- The term “systolic” was first used in this context by H.T. Kung, then at CMU; it refers to the “pumping” action of a heart.
- Machines have been constructed based on this principle, notably the iWARP, fabricated by Intel in the 80’s.

Early Systolic Operation

- Colossus Mark 2 at Bletchley Park, 1944 (http://en.wikipedia.org/wiki/Colossus_computer)
- “Colossus included the first ever use of shift registers and systolic arrays, enabling five simultaneous tests, each involving up to 100 Boolean calculations, on each of the five channels on the punched tape (although in normal operation only one or two channels were examined in any run).”

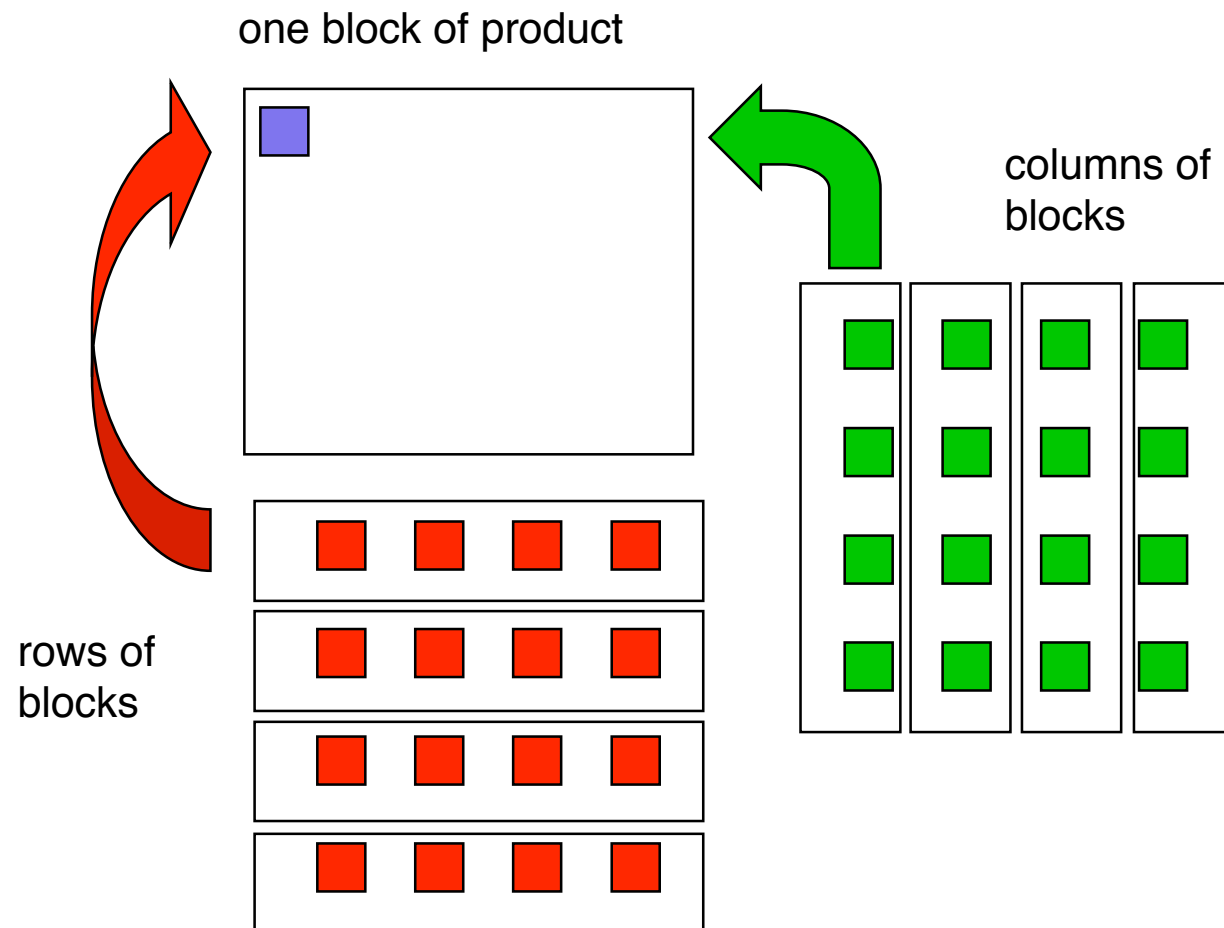
In 1994, a team led by Tony Sale (right) began a reconstruction of a Colossus at Bletchley Park. Here, in 2006, Sale supervises the breaking of an enciphered message with the completed machine.



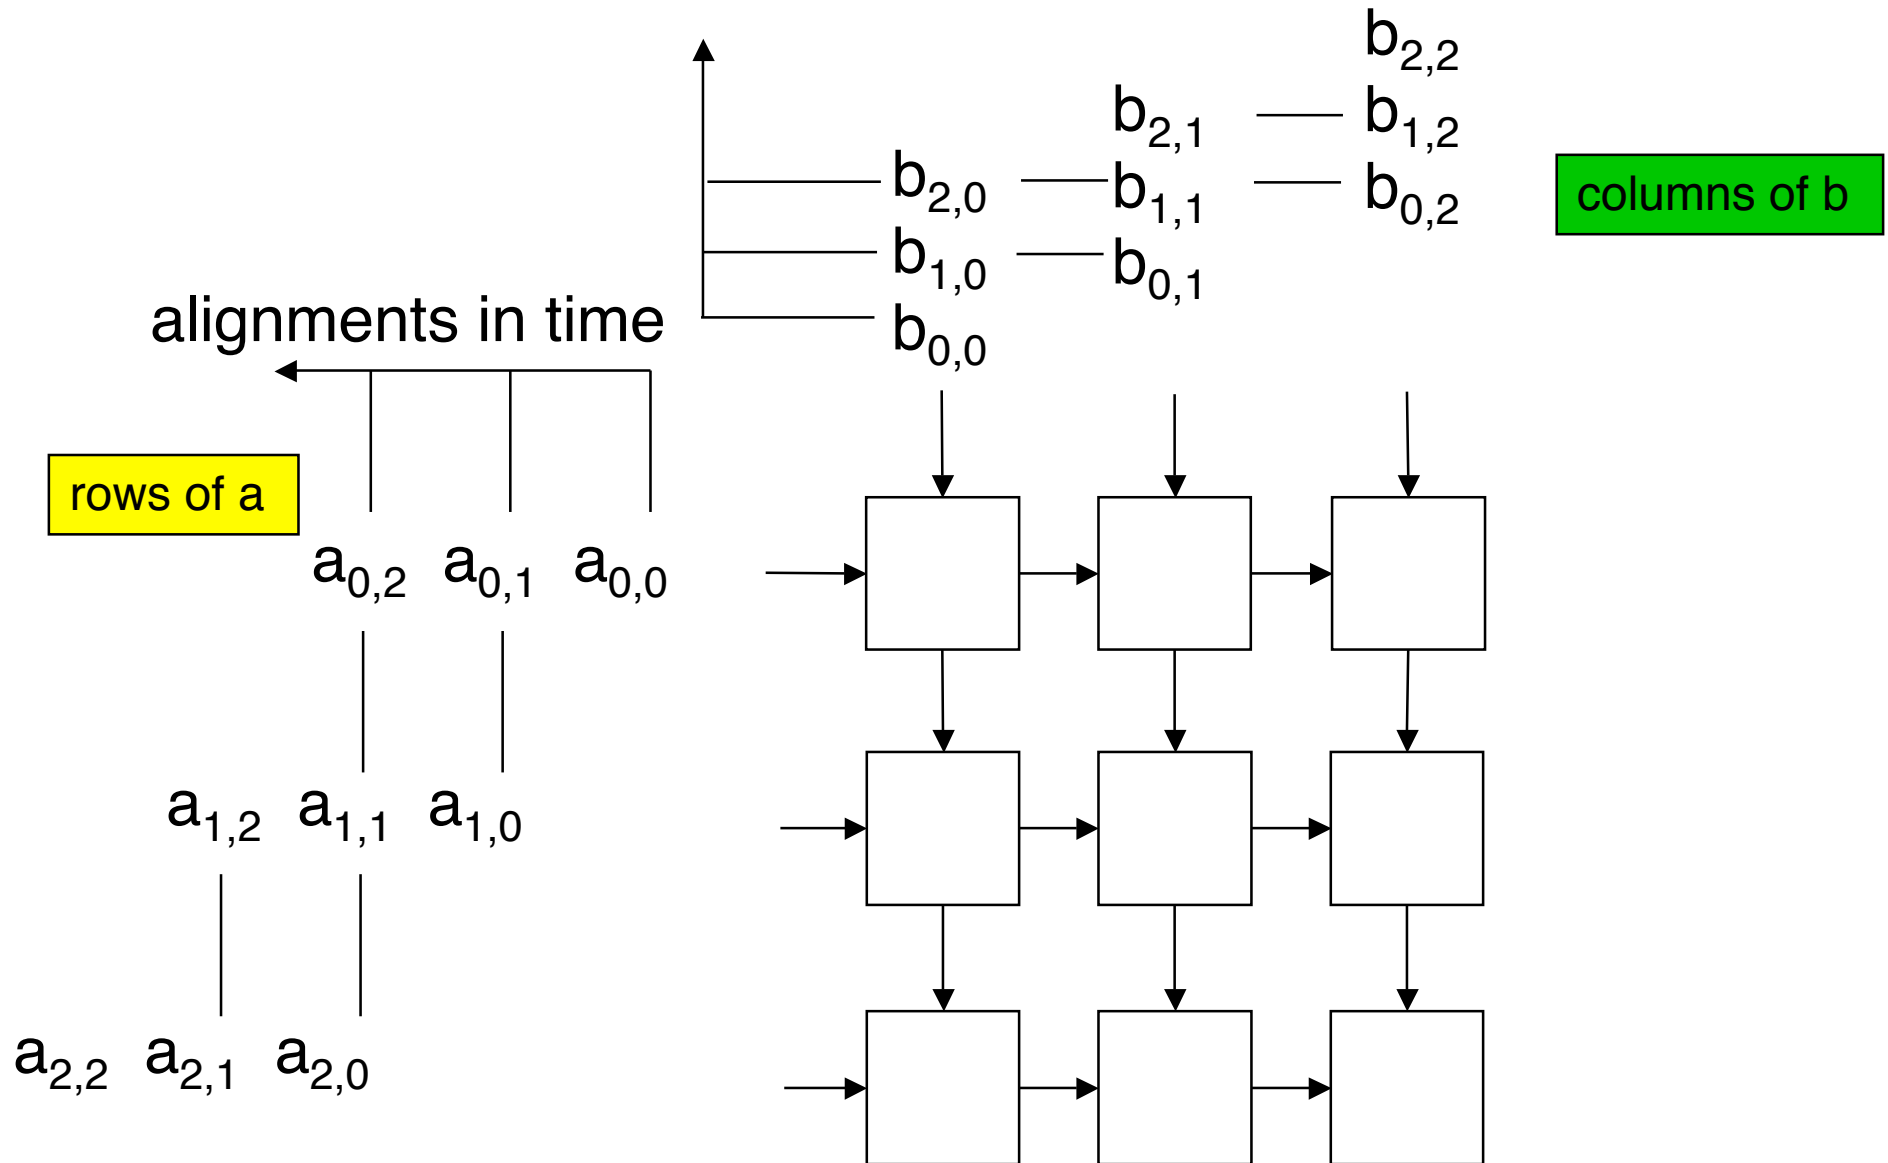
Systolic Matrix Multiplication

- Processors are arranged in a 2-D grid.
- Each processor accumulates one element of the product.
- The elements of the matrices to be multiplied are “pumped through” the array.
- We illustrate for individual elements, but can be applied to **block decompositions**.

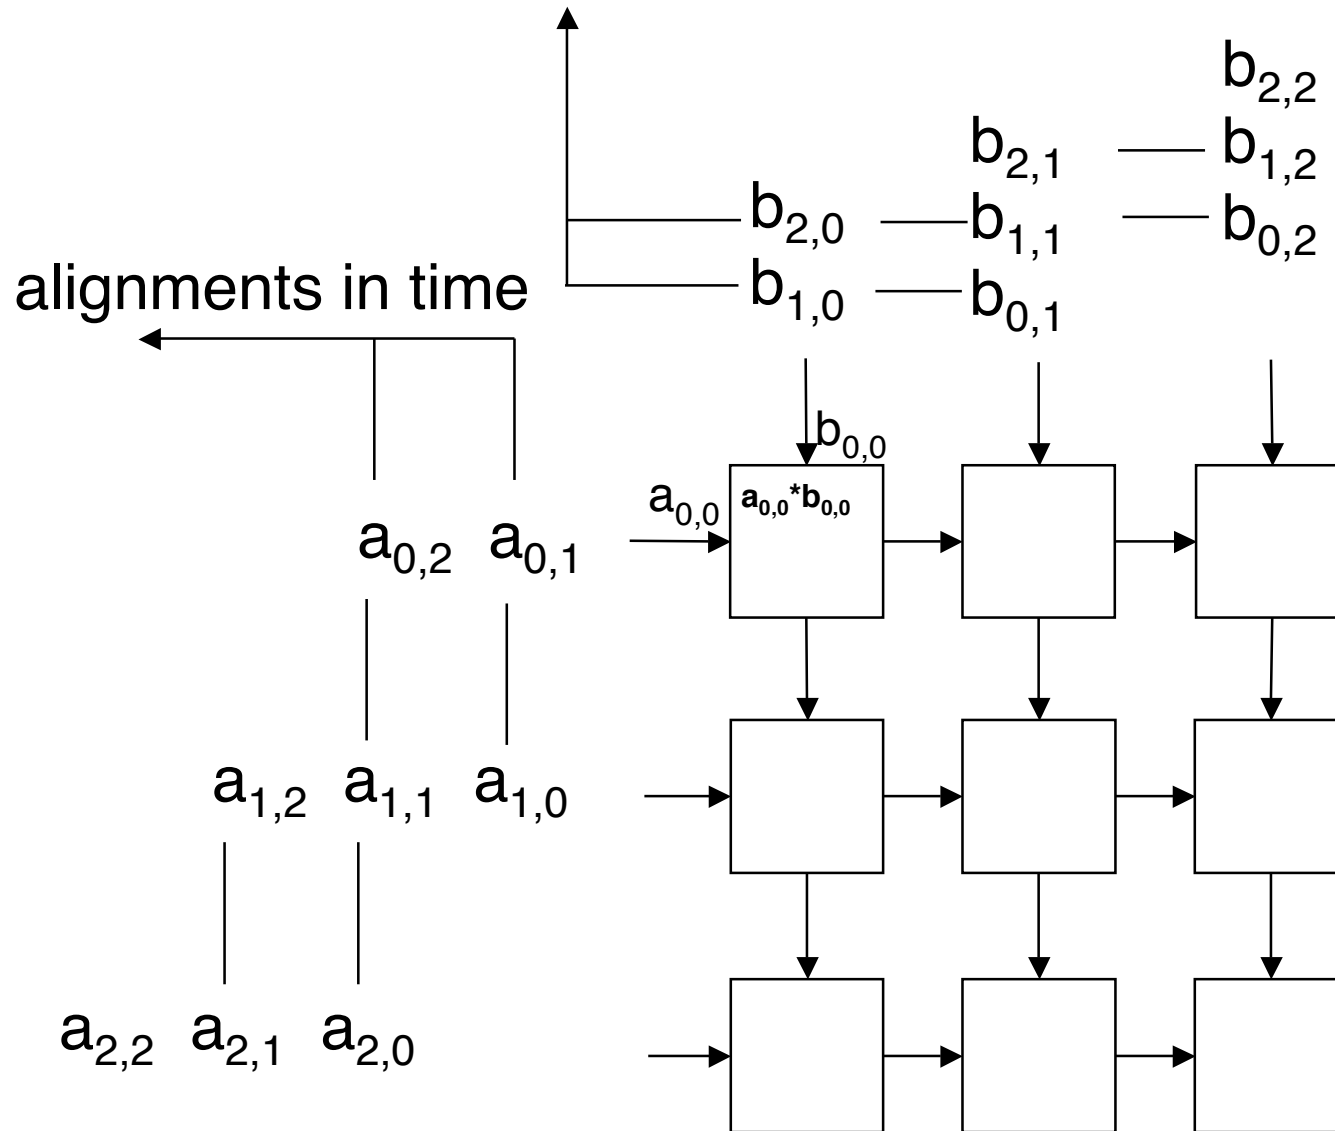
Block Multiplication



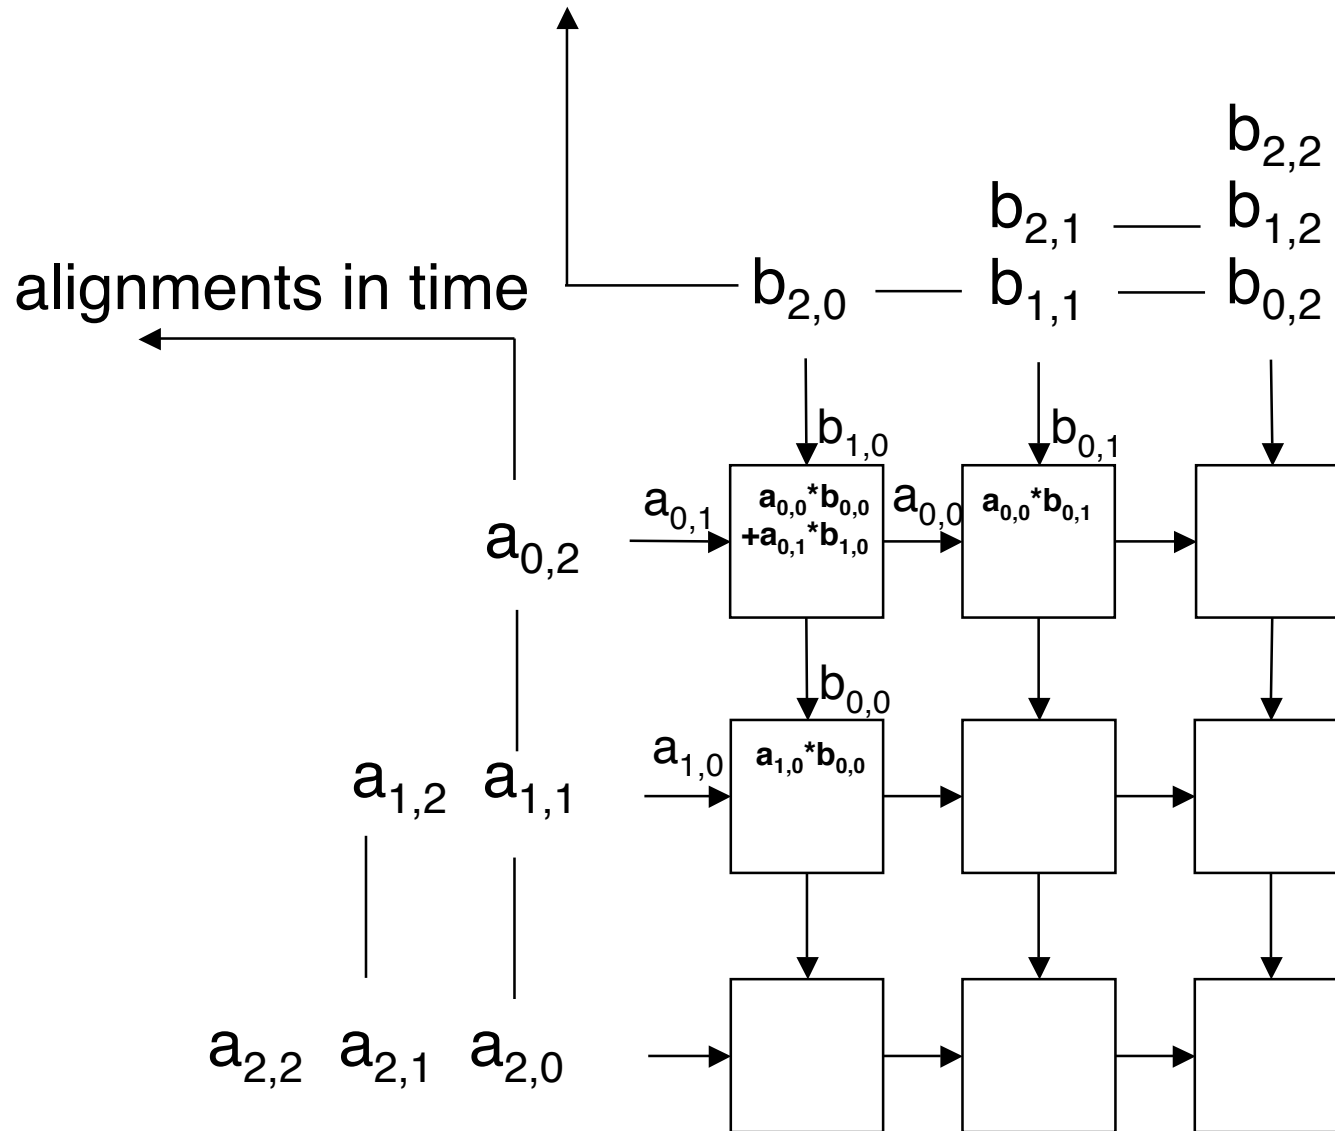
Systolic Matrix Multiplication Illustrated with two 3x3 matrices



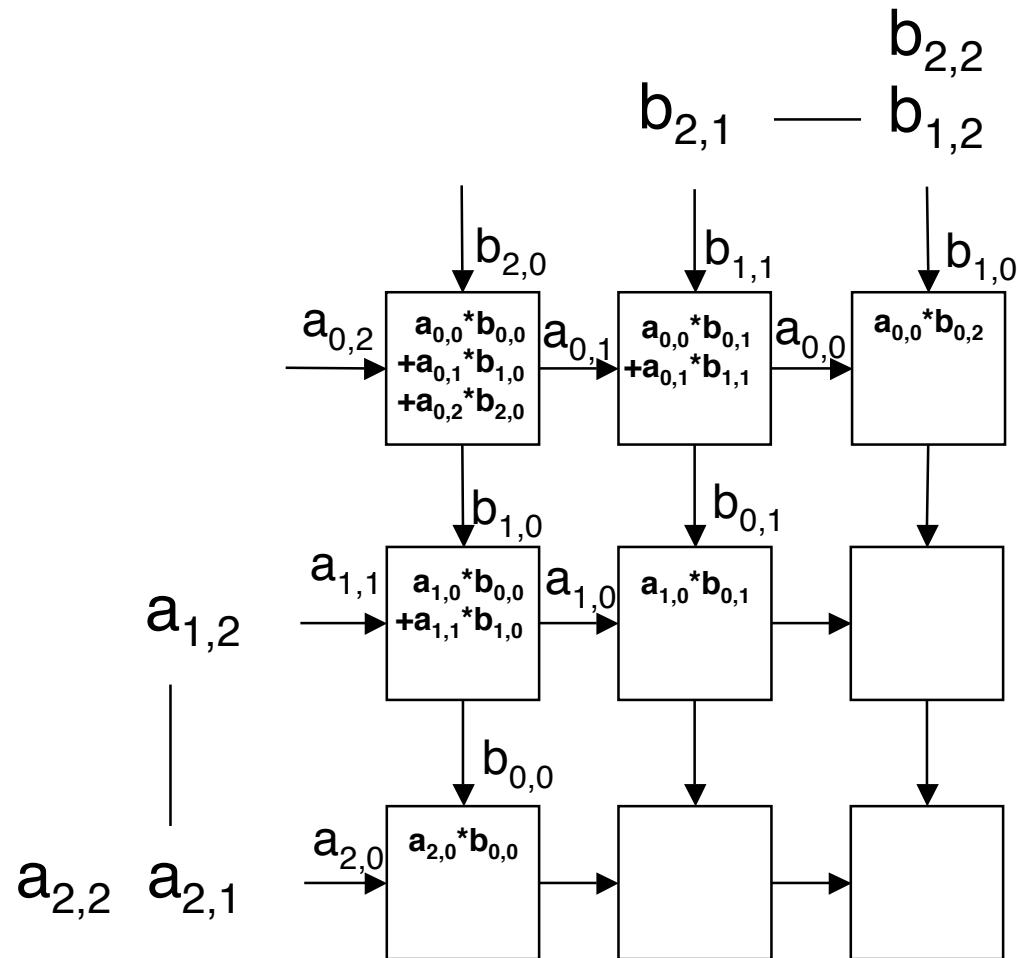
Systolic Matrix Multiplication Illustrated with two 3x3 matrices



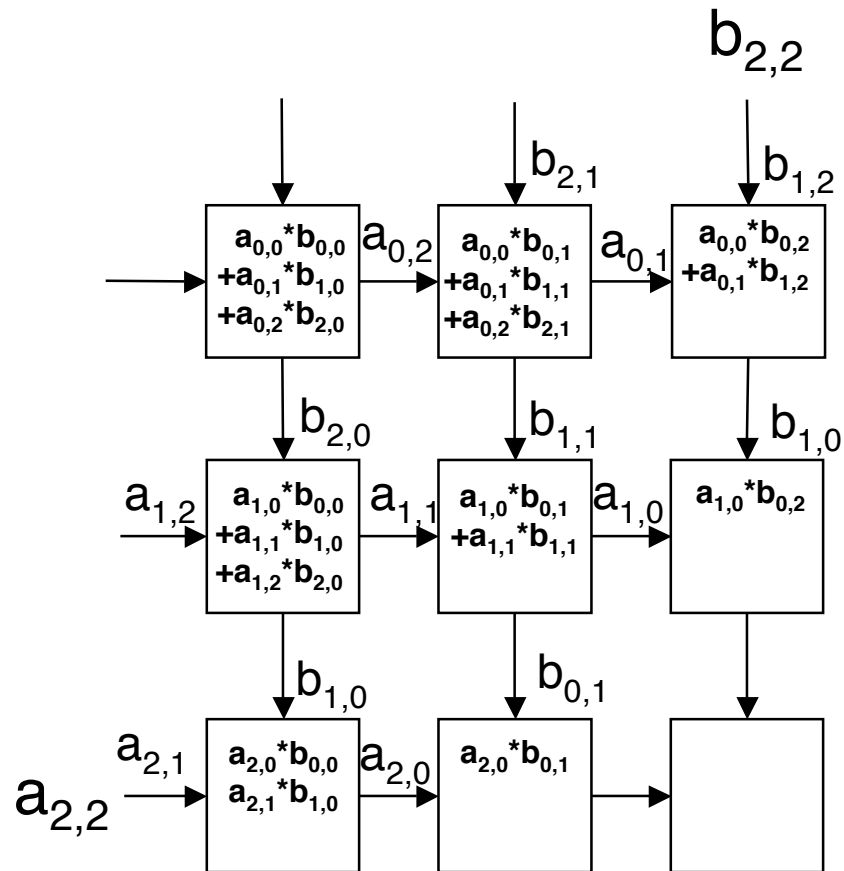
Systolic Matrix Multiplication Illustrated with two 3x3 matrices



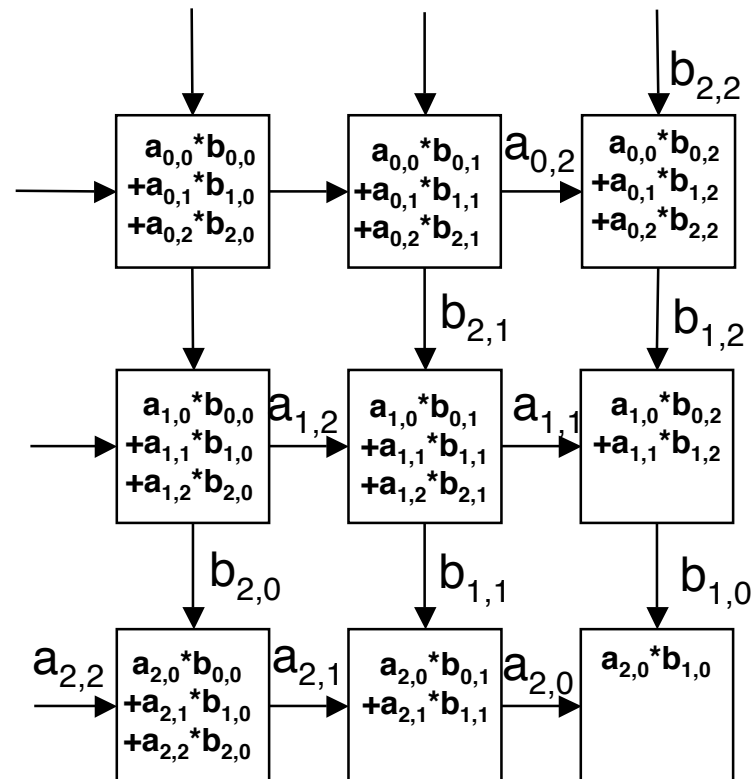
Systolic Matrix Multiplication Illustrated with two 3x3 matrices



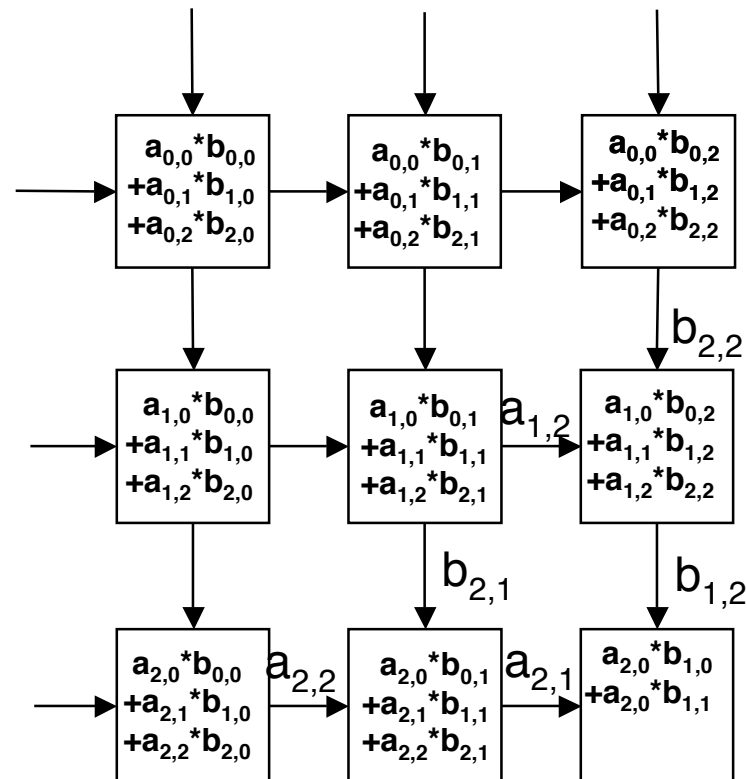
Systolic Matrix Multiplication Illustrated with two 3x3 matrices



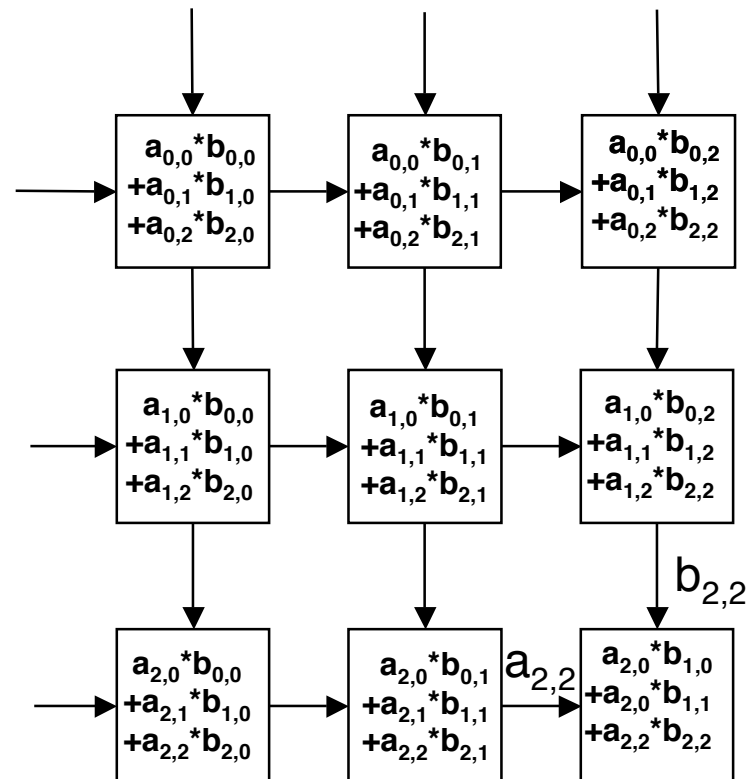
Systolic Matrix Multiplication Illustrated with two 3x3 matrices



Systolic Matrix Multiplication Illustrated with two 3x3 matrices



Systolic Matrix Multiplication Illustrated with two 3x3 matrices



Analysis

- Assuming $n \times n$ matrices (n^2 elements)
- $O(n^3)$ operations
- $O(n)$ time steps on n^2 processors

A Related Algorithm: Cannon's Method


- Let's take another view of systolic multiplication: Consider the rows and columns of the matrices to be multiplied as **strips** that slide past each other.
- The strips are first **pre-skewed** so that the correct elements are multiplied at each time step.

Cannon's Method

- Rather than have some processors idle,
- **wrap** the array rows and columns so that **every processor is doing something on each step.**
- In other words, rather than feeding in the elements, they are **rotated** around,
- starting in an initially staggered position as in the systolic model.
- We also change the order of products slightly, to make it correspond to more natural storage by rows and columns.

Cannon's Matrix Multiply ("Pipe-Roll" method)

Note that the **a diagonal** is in the left column and the **b diagonal** is in the top row.



Products computed at each step:

$$\text{Example sum: } c_{0,2} = a_{0,2} * b_{2,2} + a_{0,1} * b_{1,2} + a_{0,0} * b_{0,2}$$

Step 1

$a_{00} * b_{00}$	$a_{01} * b_{11}$	$a_{02} * b_{22}$
$a_{11} * b_{10}$	$a_{12} * b_{21}$	$a_{10} * b_{02}$
$a_{22} * b_{20}$	$a_{20} * b_{01}$	$a_{21} * b_{12}$

Step 2

$a_{02} * b_{20}$	$a_{00} * b_{01}$	$a_{01} * b_{12}$
$a_{10} * b_{00}$	$a_{11} * b_{11}$	$a_{12} * b_{22}$
$a_{21} * b_{10}$	$a_{22} * b_{21}$	$a_{20} * b_{02}$

Step 3

$a_{01} * b_{10}$	$a_{02} * b_{21}$	$a_{00} * b_{02}$
$a_{12} * b_{20}$	$a_{10} * b_{01}$	$a_{11} * b_{12}$
$a_{20} * b_{00}$	$a_{21} * b_{11}$	$a_{22} * b_{22}$

Cannon's Method for **Block** Multiplication

- At each step, entire **blocks** are transmitted down and to the left of neighboring PE's.
- Memory space is conserved.


Fox's Algorithm

- Also for block matrix multiplication, it has a resemblance to Cannon's algorithm.
- The difference is that **on each cycle**:
 - A row block is **broadcast** to **each other** processor in the row.
 - The column blocks are rolled cyclically.

Fox's Algorithm

Step 1

$a_{00} * b_{00}$	$a_{00} * b_{01}$	$a_{00} * b_{02}$
$a_{11} * b_{10}$	$a_{11} * b_{11}$	$a_{11} * b_{12}$
$a_{22} * b_{20}$	$a_{22} * b_{21}$	$a_{22} * b_{22}$



A different **row element** of **a** is **broadcast** in each step.

Step 2

$a_{01} * b_{10}$	$a_{01} * b_{11}$	$a_{01} * b_{12}$
$a_{12} * b_{20}$	$a_{12} * b_{21}$	$a_{12} * b_{22}$
$a_{20} * b_{00}$	$a_{20} * b_{01}$	$a_{20} * b_{02}$



b column elements are **rolled** each step

Step 3

$a_{02} * b_{20}$	$a_{02} * b_{21}$	$a_{02} * b_{22}$
$a_{10} * b_{00}$	$a_{10} * b_{01}$	$a_{10} * b_{02}$
$a_{21} * b_{10}$	$a_{21} * b_{11}$	$a_{21} * b_{12}$

Fox vs. Cannon

- Fox does not require **pre-skewing**.
- Broadcast *may* be cheaper than multiple block transfers.

Other Uses of Systolic Arrays

- Cholesky Factorization of Matrices
- Digital filters (FIR)
- Genome Matching (Smith-Waterman Algorithm, Paracel Inc.)
- <http://www.timelogic.com/technology.html>
(FPGA-based computers)

References

- *Ziad A.A. Alqadi, Musbah Aqel, and Ibrahiem M. M. El Emary, Performance Analysis and Evaluation of Parallel Matrix Multiplication Algorithms, World Applied Sciences Journal 5 (2): 211-214, 2008*
- Analysis of a Class of Parallel Matrix Multiplication Algorithms
<http://www.cs.utexas.edu/users/plapack/papers/ipps98/ipps98.html>
- Two-Dimensional **Block Cyclic** Data Distribution as a Key to Load Balancing and Software Reuse (part of ScaLAPACK User's Guide)
<http://www.netlib.org/scalapack/slug/node110.html>