

# Reconfigurable Computing

## Customisation

*“The only man who behaves sensibly is my tailor; he takes my measurements anew every time he sees me, while all the rest go on with their old measurements and expect me to fit them”*

– George Bernard Shaw

Philip Leong (philip.leong@sydney.edu.au)  
School of Electrical and Information Engineering

<http://phwl.org/talks>



Permission to use figures have been gained where possible. Please contact me if you believe anything within infringes on copyright.

- › The more you know about a problem, the better you can customise a solution
- › It is hard to find general principles to guide opportunities for customisation. We will use 2 examples
  - Distributed arithmetic (where we rearrange equations for doing a dot product)
  - DNA sequence matching (where we fit a problem requiring 4 slices into 3)

# Distributed Arithmetic



- › Calculate complex function  $f(x)$
- › If number of bits of  $x$  is small, just make a ROM-based lookup table
  - $F[x]$  where we precompute for all possible  $x$

# Distributed Arithmetic – Reorganising Equations

› Consider a dot product  $\langle \mathbf{w}, \mathbf{x} \rangle = \sum_{k=0}^{K-1} w[k]x[k]$

where the  $\mathbf{w}$  are fixed coefficients and  $x[k]$  are 2's complement fractions

› Recall if  $x[k] = \{b_{N-1}[k], \dots, b_0[k]\}$

$$x[k] = 2^{-F} (-b_{N-1}[k]2^{N-1} + \sum_{n=0}^{N-2} b_n[k]2^n),$$

then

$$\langle \mathbf{w}, \mathbf{x} \rangle = \sum_{k=0}^{K-1} w[k] (-b_{N-1}[k]2^{N-1-F} + \sum_{n=0}^{N-2} b_n[k]2^{n-F})$$

rearranging

$$\langle \mathbf{w}, \mathbf{x} \rangle = \sum_{n=0}^{N-2} \sum_{k=0}^{K-1} w[k] b_n[k] 2^{n-F} - \sum_{k=0}^{K-1} w[k] b_{N-1}[k] 2^{N-1-F}$$

$$\langle \mathbf{w}, \mathbf{x} \rangle = \sum_{n=0}^{N-2} \sum_{k=0}^{K-1} w[k] b_n[k] 2^{n-F} - \sum_{k=0}^{K-1} w[k] b_{N-1}[k] 2^{N-1-F} \quad (1)$$

Let's assume  $F=N-1$  (this means our implementation will only shift in 1 direction)

Also remember  $b_n[k]$  is binary and  $w[k]$  are constants so let's store

$$m[a_n] = \sum_{k=0}^{K-1} w[k] b_n[k],$$

where  $a_n = \{b_n[K-1], \dots, b_n[0]\}$

Then (1) becomes

$$\begin{aligned} \langle \mathbf{w}, \mathbf{x} \rangle &= \sum_{n=0}^{N-2} m[a_n] 2^{n-N+1} - m[a_{N-1}] 2^0 \\ &= m[a_0] 2^{1-N} + \dots + m[a_{N-2}] 2^{-1} - m[a_{N-1}] 2^0 \end{aligned}$$

- ›  $\langle w, x \rangle$  can be computed using the recurrence

$$s_0 = m[a_0]$$

$$s_i = s_{i-1}2^{-1} + m[a_i] \quad (i = 1, 2, \dots, N - 2)$$

$$s_i = s_{i-1}2^{-1} - m[a_i] \quad (i = N - 1)$$


- › Can you draw a datapath which can implement this?

>  $\langle w, x \rangle$  can be computed using the recurrence

$$s_0 = m[a_0]$$

$$s_i = s_{i-1}2^{-1} + m[a_i] \quad (i = 1, 2, \dots, N - 2)$$

$$s_i = s_{i-1}2^{-1} - m[a_i] \quad (i = N - 1)$$





Suppose  $\mathbf{w} = [w[1], w[0]] = [0.375, -0.5]$  and  $\mathbf{x} = [-0.75, 0.125]$ ,  $N = 6$  and  $K = 2$  is the length of the vectors. In binary,  $\mathbf{x} = [1.0100, 0.0010]$ . Thus

$$\langle \mathbf{w}, \mathbf{x} \rangle = w[1] \times x[1] + w[0] \times x[0] \quad (26)$$

$$= 0.375 \times -0.75 + -0.5 \times 0.125 \quad (27)$$

$$= 0.375 \times (-2^0 + 2^{-2}) + -0.5 \times (2^{-3}) \quad (28)$$

Distributed arithmetic rearranges this computation to be

$$\begin{aligned} \langle \mathbf{w}, \mathbf{x} \rangle &= 2^{-3} \times -0.5 + \\ &= 2^{-2} \times 0.375 - \\ &= 2^0 \times 0.375. \end{aligned}$$

$a_1$	$a_0$	$m[a_1, a_0]$	Numeric Value
0	0	0	0
0	1	$w[0]$	-0.5
1	0	$w[1]$	0.375
1	1	$w[0] + w[1]$	-0.125

$a$	$m[a]$	$\pm$	$s_i$
00	0	+	$s_0 = 0$
00	0	+	$s_1 = 0$
01	-0.5	+	$s_2 = -0.5$
10	0.375	+	$s_3 = 0.125$
00	0	+	$s_4 = 0.0625$
10	0.375	-	$s_5 = -0.34375$


- › Dot products can be computed in time independent of  $K$  if it is not too big
- › Precomputation often a useful customisation trick
  - Always think about whether lookup tables can be applied if the number of input bits is small
  - c.f. CORDIC which also uses a small lookup table and rearrangement of computation

- › Stanley A. White, “Applications of Distributed Arithmetic to Digital Signal Processing: A Tutorial Review”, IEEE ASSP Magazine, July 1989

# DNA Sequence Matching

- › Background
- › Smith Waterman Algorithm
- › Improved cell
- › Results
- › Conclusion


- › DNA chain – long, unbranched polymer
- › Basic unit of DNA – Nucleotides
  
- › Nucleotides –  
4 different bases –  
adenine (A), cytosine (C),  
guanine (G) and thymine (T)
- › Long sequence of A, C, G or T nucleotides



- › The sequence - one of the strands of DNA running in a particular direction

```

GCCGCACGNTCTCATCGGTCCCGGCCGCCATGGCG
GCAGGAATTTCGATTCTATAGCANCCAGCAGACCACCT
ACGGCATATGGCAGCACAGGAAATGCTCTCTGCCAGC
ATCTCCANACCCGCATCCTGCAGACCTGCAGCGTG
CCTCATGCCAACATGGTCAACGGNGCCAACACTCACTGC
AAGGAGCTCTGGCTCCACGCCTCTATAAGTTCCCTGA
GCATGGTTTGGGTGGGGGCTCTTGTGCTTTGACCCA
CAGCTTCCCGCCGCTGCCCCAGGCCCTGCTCACGGA
CGAACCCACTCTAGGTGACATCAAGCAGGAGCTGCG
CAGGAAAAGTCGGCCCCTGGAAGAGCCGCCCGATAT
GGACTCACCTCAGATCCGTGAACTGGAGAAGTTTGCT
AATGACTTCAAACCTGAGGAGGATCAAAA
    
```




- › Find edit distance between two strings (Similarity between two sequences)
- › Identify the protein encoded by the unknown DNA sequence
- › Decode the relation between disease and inheritance
- › And others...



› Most popular: BLAST, FASTA

- Heuristic method

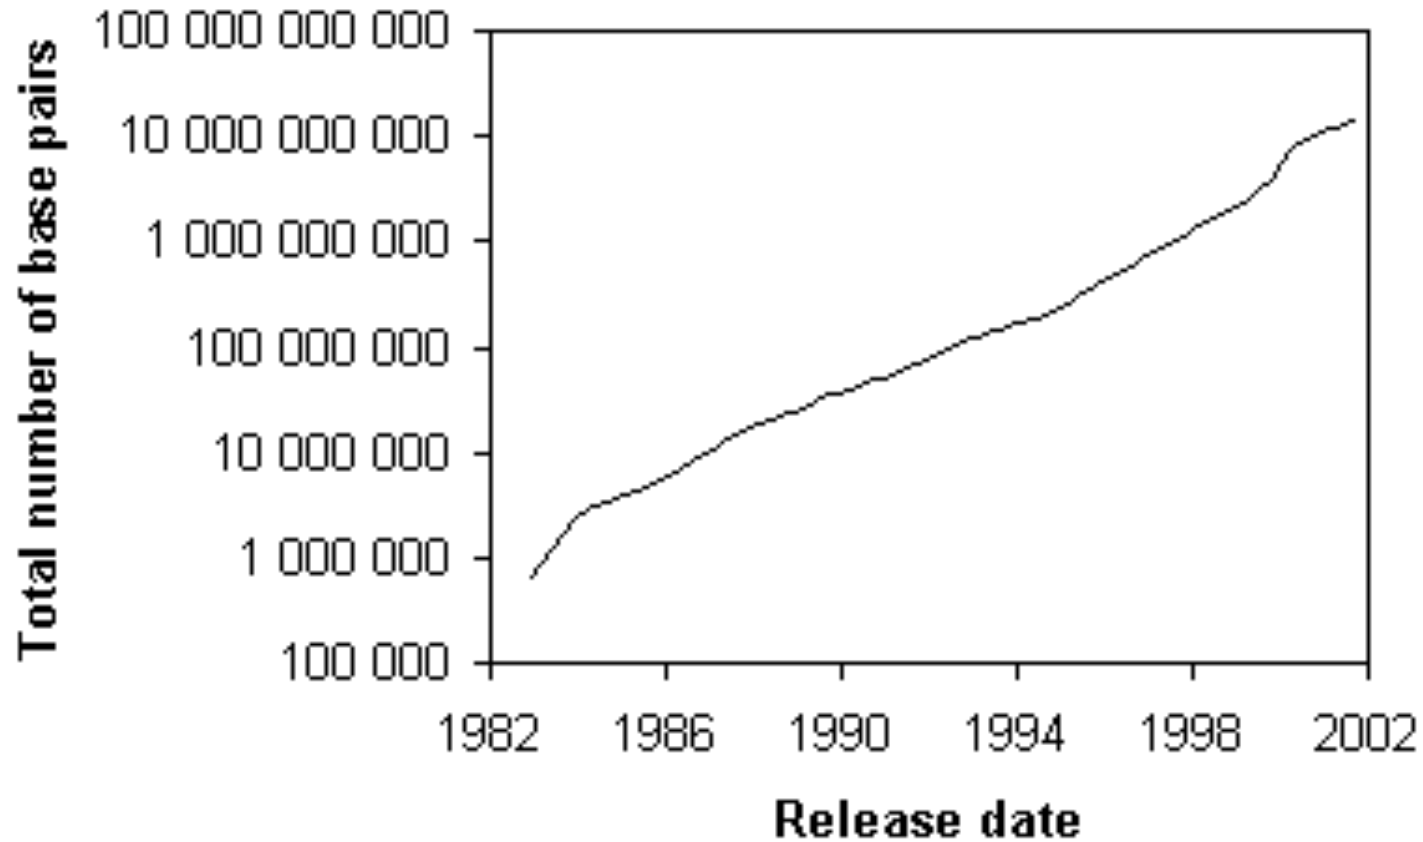


- Fast, moderate precision

› Smith-Watermann Algorithm

- Higher time complexity
- Approx  $\frac{1}{2}$  errors/query of BLAST/FASTA

## Public databases of DNA sequence of species



The graph shows the growth of GenBank

- › Splash 2 system, Hoang - 1993
  - 33 LUT/FF pairs per cell (XCV4010)
- › Xilinx, Guccione and Keller - 2002
  - Optimized Smith Waterman core uses **runtime reconfiguration** to store one of the strings and to implement the insertion, substitution and deletion penalties
  - 6 LUT/FF pairs per cell (XCV1000-6)
- › HokieGene, Virginia Tech – 2002
  - Used Xilinx's core on an Osiris board (XCV6000-4)

		T <sub>0</sub>		
		A	C	G
S <sub>0</sub>	A	<del>0</del>	<del>1</del>	<del>2</del>
	T	2	1	2
	C	3	2	1
	G	3	2	<b>2</b>

$$d = \min \begin{cases} a & \text{if } S_i = T_j \\ a + 2 & \text{if } S_i \neq T_j \\ b + 1 \\ c + 1 \end{cases}$$



**Alignment Score of "ACG" and "ATC"**




		A	C	T	A	T	C	G	T	A
	0	1	2	3	4	5	6	7	8	9
A	1	0	1	2	3	4	5	6	7	8
G	2	1	2	3	4	5	6	5	6	7
T	3	2	3	2	3	4	5	6	5	6
A	4	3	4	3	2	3	4	5	6	5
C	5	4	3	4	3	4	3	4	5	6
G	6	5	4	5	4	5	4	3	4	5
T	7	6	5	4	5	4	5	4	3	4
G	8	7	6	5	6	5	6	5	4	5

```
ACTATCGTA
| | | | |
AGTA_CGTG
```

## Simplified Equation (Lipton and Lopresti)

$$d = \begin{cases} a & \text{if } ((b \text{ or } c) = a - 1) \text{ or } (S_i = T_j) \\ a + 2 & \text{if } ((b \text{ and } c) = a + 1) \text{ and } (S_i \neq T_j) \end{cases}$$

- › d can only take two values: a or a+2
- › Change in d encoded in 1-bit (accumulated with a counter at end of systolic array)

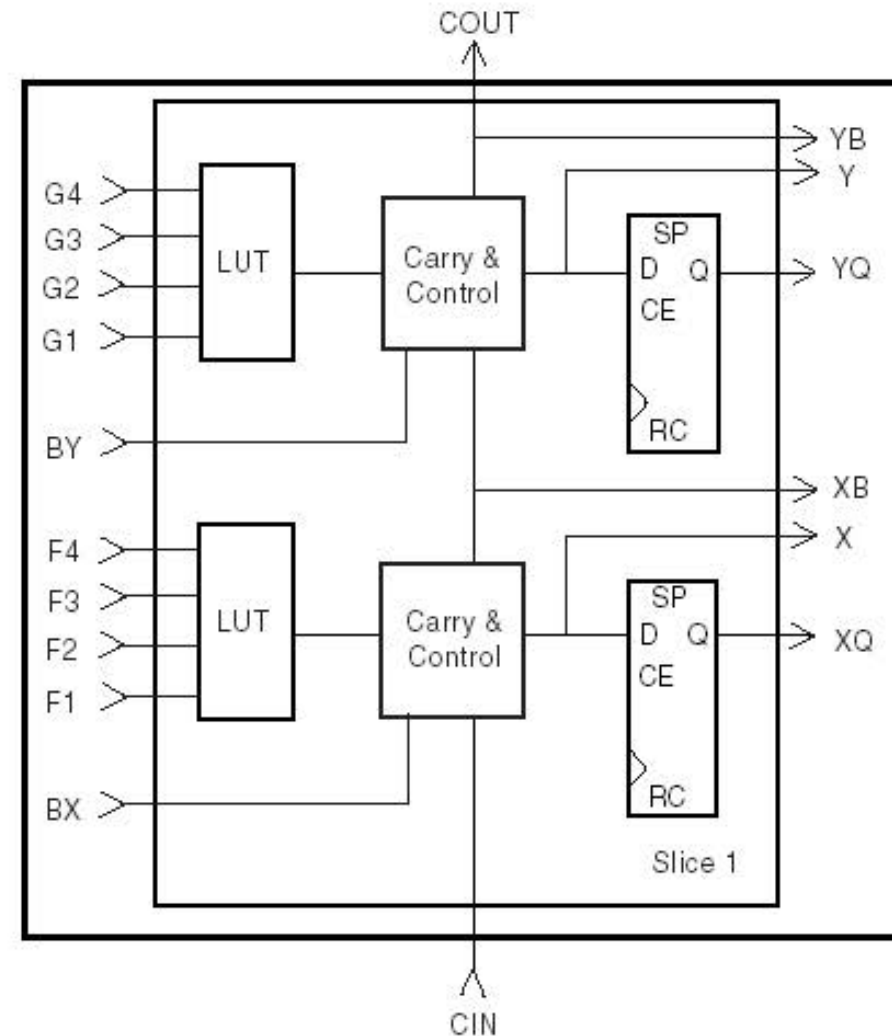


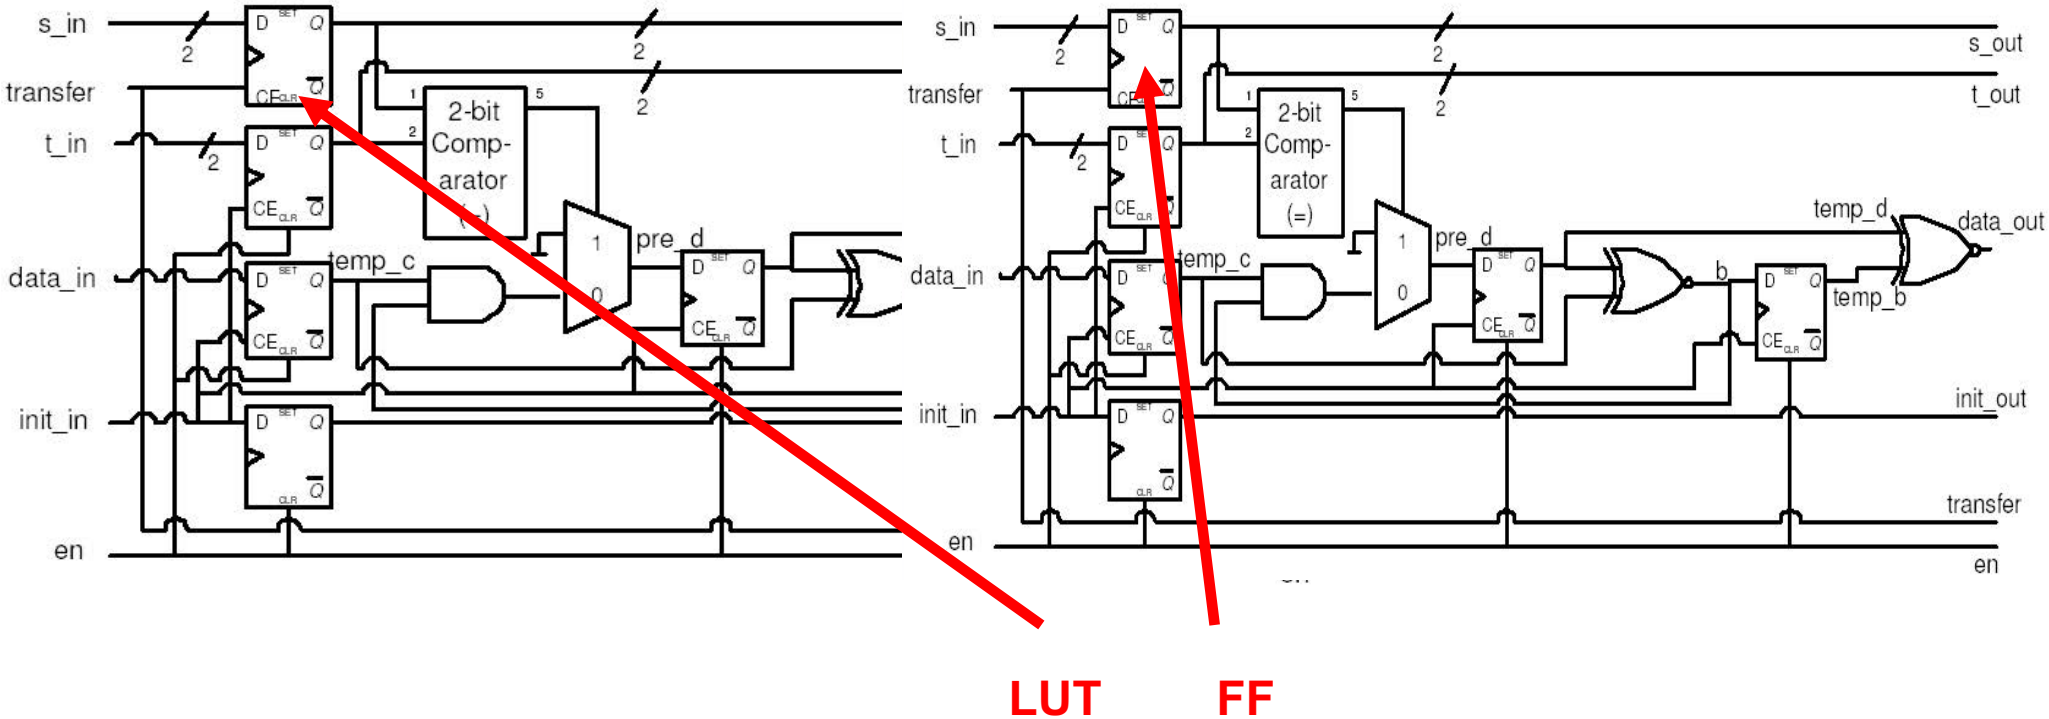
4 LUTs and 8 FFs = 4 slices

Guccione's cell uses 3 slices by writing string in bitstream



- › Cell uses 4 LUTs, 8 FFs
- › LUTs can serve as FFs using the distributed RAM feature
- › Want to use 2 LUTs to implement 2 FFs
- › Problem: there is no input which connects directly to the FF (can only do it via the LUT)





Merge two cells so that FFs are connected as in the slice  
(8 LUTs, 4 RAMs, 12 FFs) = 6 slices (2 cells)

Sequence (S)

e.g. "ACGT"

Changing Sequence (T)


e.g. "ATCC"

"ATGC"

"GCTC"


"CGGT"

.....



- › Input the S sequence
- › Stream the changing sequences T through the array
- › Read back the “alignment score”

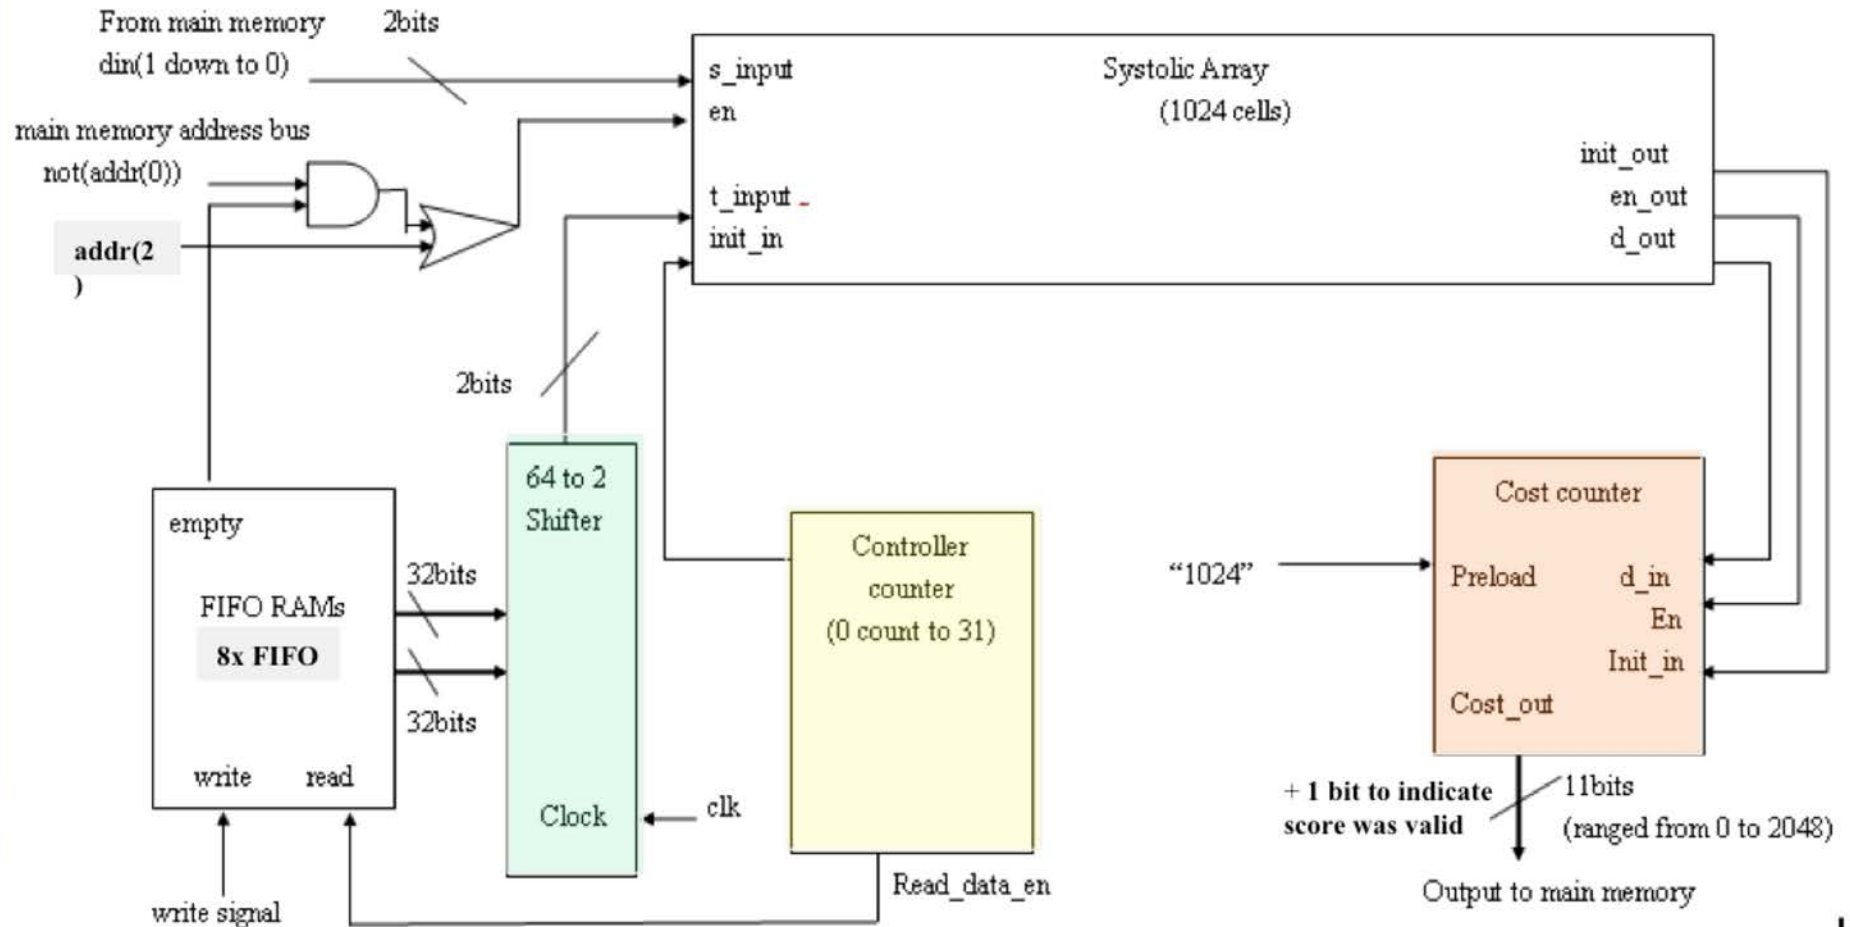





		$S_i$	$S_j$	$S_k$
		A	C	G
		0	1	2
$T_m$	A	1	0	2
$T_n$	T	2	1	3
	C	3	2	2

Red arrows point from the circled '0' and '1' in the table to the  $D_j$  input of the second and third blocks respectively. A green arrow points from the circled '1' in the table to the  $T_i$  input of the second block.

and so on...




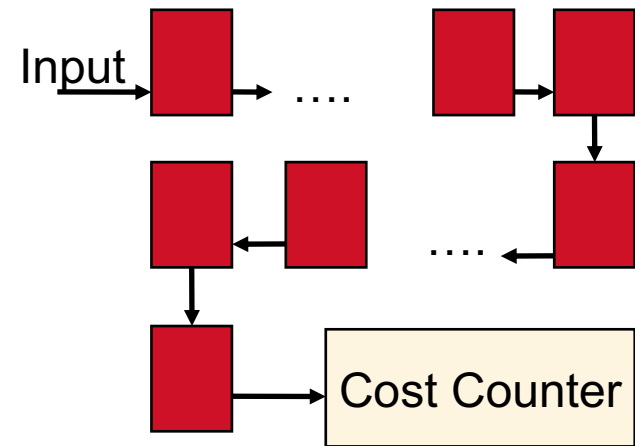
Compare 2 sequences



FIPOs, shifter,  
controller put here  
to have smaller  
delay

Systolic Array of  
2048 PE

**2048 PEs**



**4032 PEs**

- › Implemented on a Xilinx XCV1000E-6
- › 4032 cells implemented
- › Operates at 202 MHz
- › Performance of 814 billion connection updates per second
- › Correctness verified on Pilchard system



<b>System</b>	<b>Device</b>	<b>Performance (billion CUPS)</b>
<b>Splash 2</b>	<b>16× XC4010</b>	<b>2.7</b>
<b>Paracel</b>	<b>ASIC</b>	<b>1.9</b>
<b>Celera (software)</b>	<b>800 DEC Alphas</b>	<b>0.3</b>
<b>JBits</b>	<b>XCV1000-6</b>	<b>757</b>
<b>This work</b>	<b>XCV1000-6</b>	<b>742</b>

- › A new Smith-Waterman systolic cell was presented
- › Proposed approach has the same area and speed of previous work
- › Does not require runtime reconfiguration
  - Changing S is much faster
  - Can be adapted to other FPGA and ASIC families

- › C.W. Yu, K.H. Kwong, K.H. Lee, and P.H.W. Leong. A Smith-Waterman systolic cell. In Proc. International Conference on Field Programmable Logic and Applications (FPL), pages 375–384, 2003