Scalar product

$$\mathbf{a} \cdot \mathbf{b} = \langle \mathbf{a}, \mathbf{b} \rangle = \sum_{i=1}^{n} a_i b_i = a_1 b_1 + a_2 b_2 + \cdots + a_n b_n$$

- Convolution
- FIR filter (direct)
- IIR filter (direct - use for low orders only!!)

$$w(n) = x(n) - \langle a[k], w[n-k] \rangle$$

$$y(n) = b_0 w(n) + \langle b[k], w[n-k] \rangle$$

- matrix multiply $y = a \times b \longrightarrow y_{km} = \langle a_{k,:}, b_{:,m} \rangle$
- Operation: repeat (length times)
  - fetch a
 - fetch b
  - multiply and accumulate

Digital Signal Processors (DSP)

- Specialization for scalar product (and FFT)
- Single-cycle processing (memory throughput):
  - parallelism (pipelining)
  - Harvard architecture (program (P), data (X), data (Y) memories)
- Desing for embedding – I/O & host interfaces etc.
- Special arithmetic modes: rounding, saturation
- ALU in fraction mode
- ALU overflow space
- butterfly implementation (+ and -)
- Special addressing modes:
  - circular buffer,
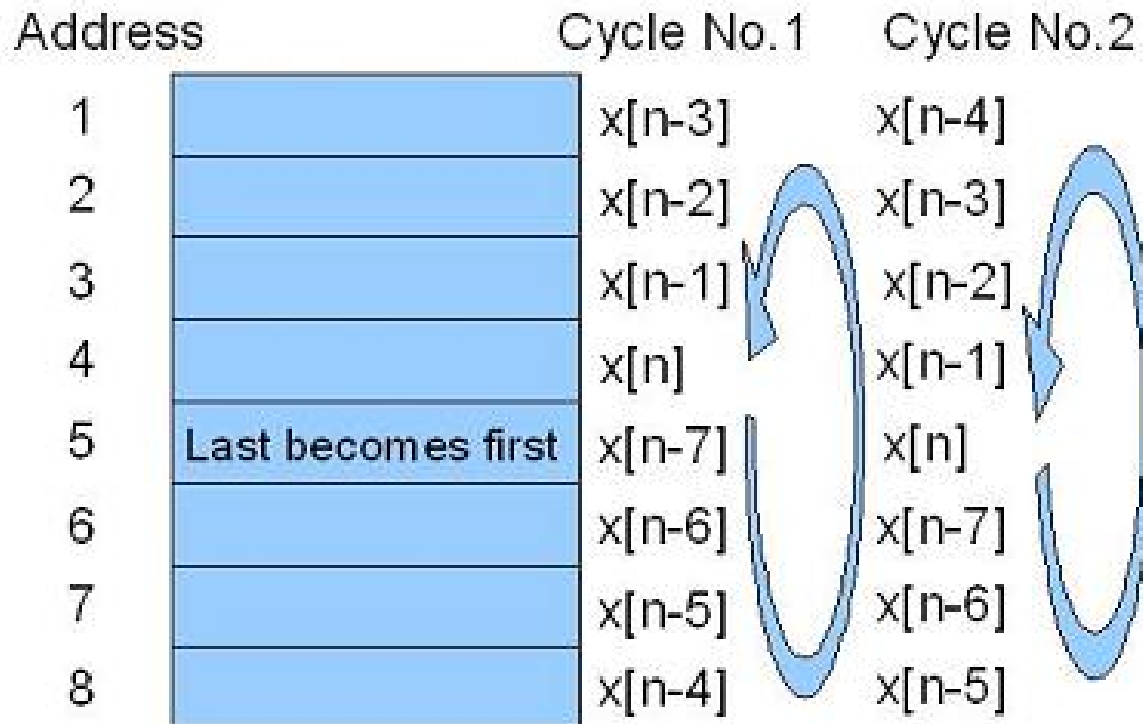 - bit-reversal,
  - matrix address

- Hardware speedups
  - Hardware loop
  - Hardware stack
  - Fast interrupts
  - Single-cycle instructions
  - RISC..

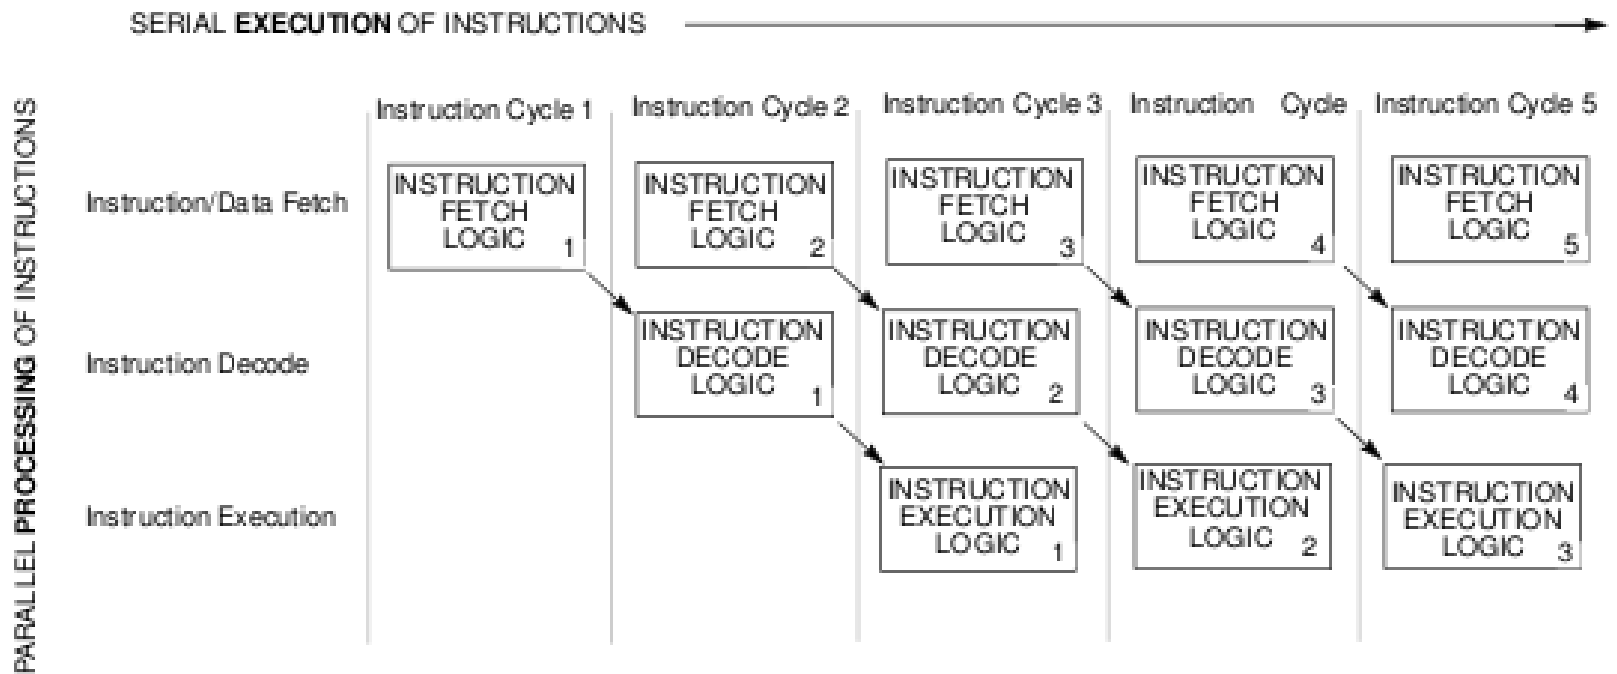# Serial I/O

Galvanic isolation:

ADuM1100

Circular buffer

Instruction pipeline

SERIAL **EXECUTION** OF INSTRUCTIONS

Manufacturers

- Texas Instruments (TMS family)
  - From simple to multiprocessor
  - fixed point and floating point
- Analog Devices (ADSP family, codename "SHARC" etc.)
- Motorola $\longrightarrow$ now Freescale Semiconductors (DSP56K family)
- vector units of general-purpose μP
- graphic processors: IBM Cell BE, Nvidia CUDA

Also "DSP cores" - VHDL or silicon IP blocks to be embedded into VLSIs.

DSP56002

- Arithmetics:
  - 24-bit (144 dB), fixed-point fraction (-1.0 till +1.0)
  - 2x24+8 bits in accumulators
  - rounding and saturation
- Addressing
  - indirect: `X:(Rn)+Nn` (post-modification)
  - indexed: `X:(Rn+Nn)`
  - absolute: `X:7`
  - immediate: `#0.125`
  - modulus register: M=-1 (no modif) M=0 (bit-reversal) M=other (circ.buffer)
- Memory: 256(X)+256(Y), extendable to 64k

Assembly programming of DSP56002

registers	symbols			
X0, X1, Y0, Y1	x, y			
A, B	a, b	s	g, h	
R0, ..., R7	r			
N0, ..., N7	n	i		

```
abs, asl, asr, clr, neg, rnd:       abs a;
add, sub:                        add s, a;
mpy, mpyr, mac, macr: mpy ±x,y,a;
nop
```

```
move x:ea,g; from memory
move g,x:ea; to memory
move ea; (update Rn)
move g,h;
move #c,g;
```

asembly		meaning		mode
ea	X&Y	ea	R update	
(r)-n			r=r-n;	
(r)+n	(yes)		r=r+n;	
(r)-	(yes)		r=r-1;	
(r)+	(yes)		r=r+1;	
(r)	(yes)	r		indirect
(r+n)		r+n		indexed
c		c		absolute

```
macr -x0,x0,a   a,x:(r3)-   y:(r5)+n5,x0
```

Example – FIR filter

```
N        equ  8
         org  x:0
samples ds   N
         org  y:0
coeffs  dc   0.0286,0.0716,0.1683,0.2458
         org  p:64     ; start address
         init       ;
         move  #samples,r0     ;
         move  #coeffs,r4      ;
         move  #N-1,m0      ;
         move  m0,m4       ;
         repeat         ;
         in    x:(r0)
         clr   a    x:(r0)+,x0     y:(r4)+,y0      ; x(n), h(0)
         .loop   #N-1
         mac   x0,y0,a         x:(r0)+,x0     y:(r4)+,y0      ; x(n-k), h(k)
         .endl
         macr    x0,y0,a    (r0)-          ;
         out    a
         forever         ;
```

Speedups

- SIMD mode: the same operation on different data (e.g. different rows of a matrix): multiple ALU's

- Multiprocessing (with separate RAMs)

- DMA I/O

Alternatives: GPGPU

- massive multicore, massive pipelining, SIMT
- stream processing: applying the same operation ("kernel") to each element of the stream of data (vector element or vector fragment)
  a `for (i=0;i++;i<N)` loop where each i-th loop content is executed in parallel
- linear algebra libraries (CUBLAS)

Main usage: speeding up offline calculations, simulations etc., with `main()` running on a PC

For realtime work in embedded environment – DSP is more predictable!