

Mapping Signal Processing Algorithms to Architecture

Sumam David S

Head, Dept of Electronics & Communication

National Institute of Technology Karnataka, Surathkal, India

sumam@ieee.org

Objectives



■ At the end of the lecture

- have gained an overview of implementation aspects of signal processing algorithms
- have gained an understanding of the relationship between the parameters influencing the choice of implementation
- appreciate the approaches in efficient mapping signal processing algorithms to architecture
- be familiar with few transformations that help the designer to develop different solutions for a given signal processing algorithm

Organisation



- What is SP ?
- Applications of SP
- Key algorithms
- Implementation options & trade-offs
- Key approaches in mapping algorithms to architecture

SU2010-CS

Algo2Arch

3

NITK Surathkal, India



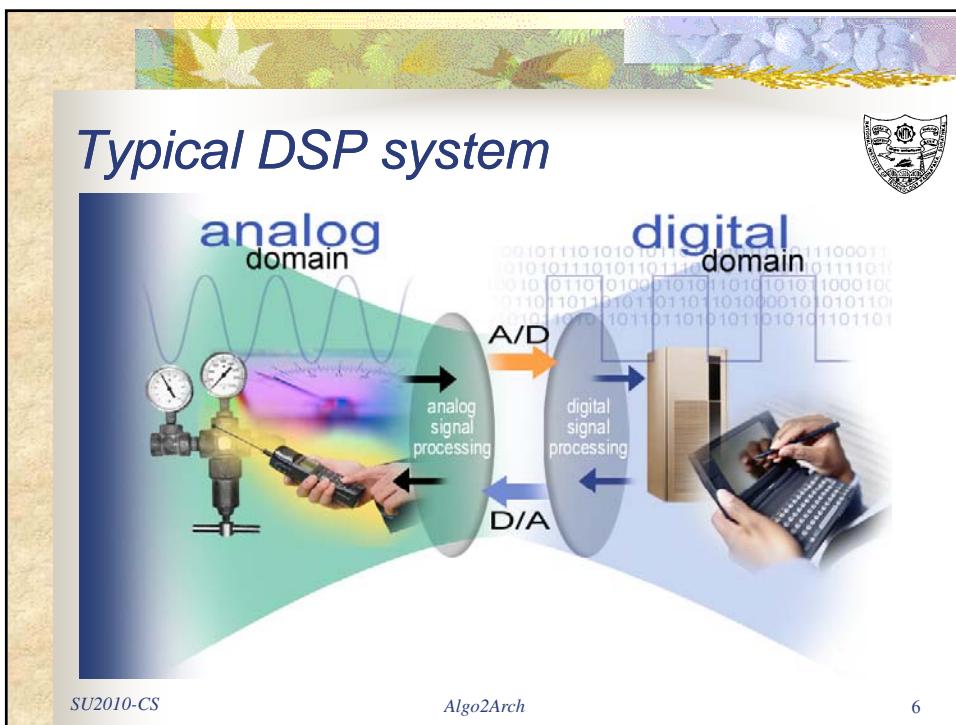
- 12°N, 74°E
- West coast of India
- Arabian Sea
- Mangalore



SU2010-CS

Algo2Arch

4

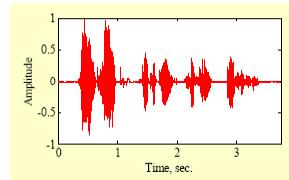




Examples of typical signals

■ Speech and music signals

- Represent air pressure as a function of time at a point in space



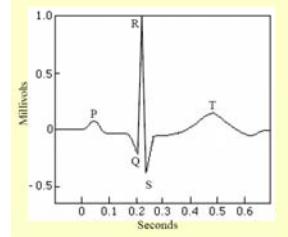
SU2010-CS

Algo2Arch

7



Biomedical signals



SU2010-CS

Algo2Arch

8

EEG

SU2010-CS

Algo2Arch

9

Seismic signal

SU2010-CS

Algo2Arch

10

Image



$I(x,y)$

SU2010-CS

Algo2Arch

11

Video



SU2010-CS

Algo2Arch

12

Digital Camera



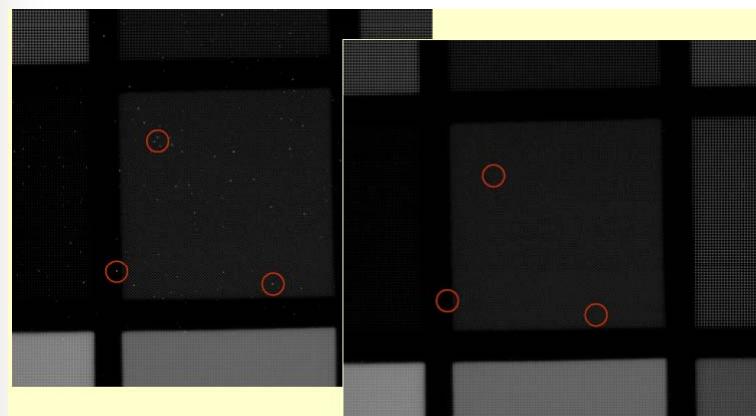
- Image Processing Algorithms
 - Bad pixel detection and masking
 - Color interpolation
 - Color balancing
 - Contrast enhancement
 - False color detection and masking
 - Image and video compression

SU2010-CS

Algo2Arch

13

Bad pixel detection & masking



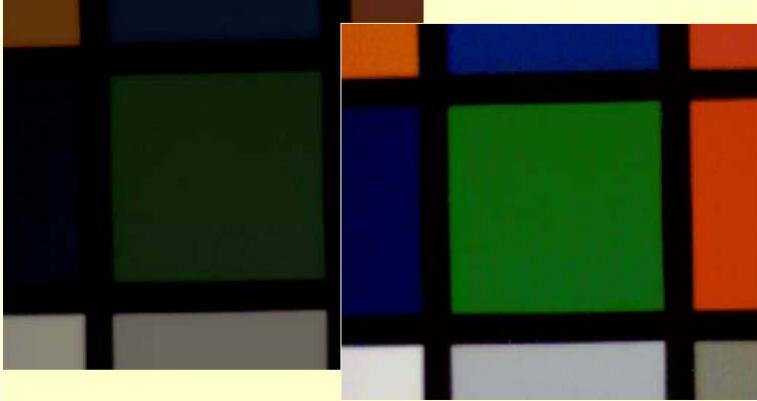
SU2010-CS

Algo2Arch

14



Color interpolation & balancing



SU2010-CS Algo2Arch 15



Digital Sound Synthesis



- Guitar with nylon strings 🎸
- Marimba 🎵
- Tenor saxophone 🎹

SU2010-CS Algo2Arch 16



Signal Compression



- Original music
 - Audio Format: PCM 16.000 kHz, 16 Bit
 - (Data size 66206 bytes) 
- Compressed music
 - Audio Format: GSM 6.10, 22.05 kHz
 - (Data size 9295 bytes) 

SU2010-CS

Algo2Arch

17



Image Compression



8 bpp

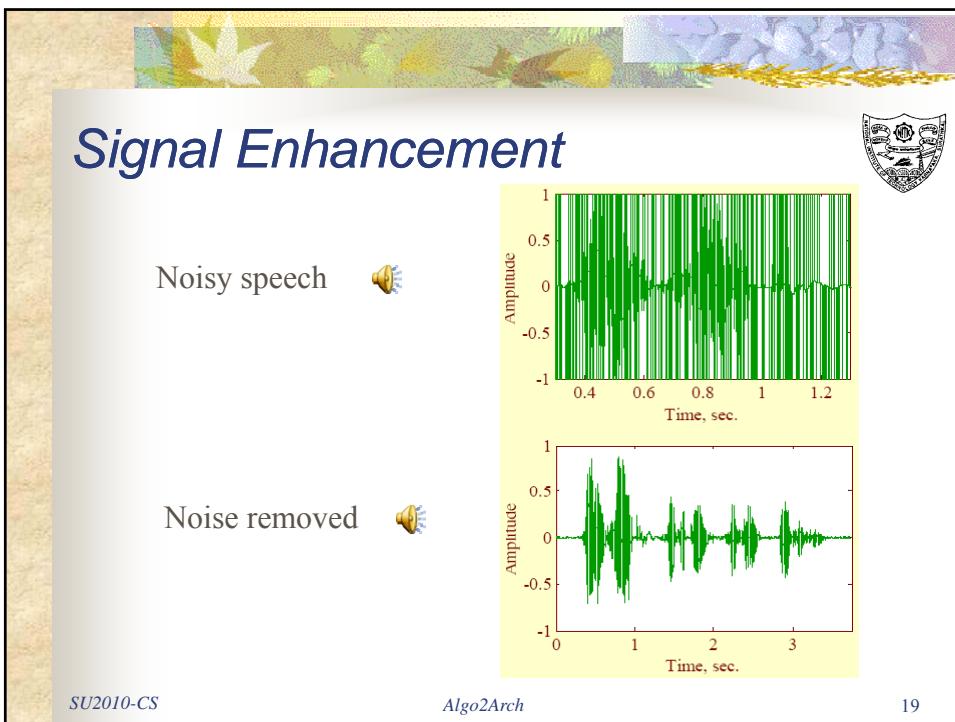


0.5 bpp

SU2010-CS

Algo2Arch

18



SU2010-CS

Algo2Arch

19



SU2010-CS

Algo2Arch

20

Noise removal

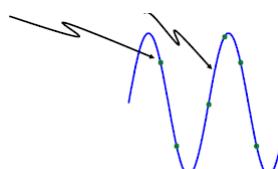


20% pixels corrupted with additive impulse noise Noise-removed version

SU2010-CS Algo2Arch 21

Signal Processing

- Works on discrete samples of a continuous time signal
- Real time requirement
- Data driven
- Programmable or custom DSPs



SU2010-CS Algo2Arch 22



Implementation choice

- Depends on

- Sampling rate
- Throughput
- Power - energy
- Area
- Wordlength – precision
- Flexibility
- Time market
- Volume

SU2010-CS

Algo2Arch

23



Examples of DSP Primitives

- Convolution
- Digital Filters
 - Finite Impulse response (FIR)
 - Infinite Impulse Response (IIR)
- Correlation
- Discrete Fourier Transform / Fast Fourier Transform (FFT)
- Discrete Cosine Transform (DCT)
- Least Mean Square (LMS)

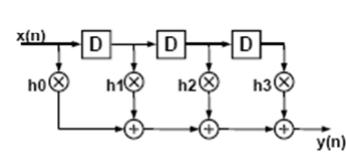
Applications often comprised of many primitives

SU2010-CS

Algo2Arch

24

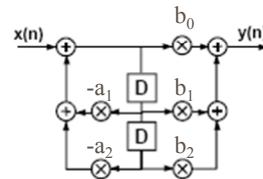
Two basic DSP structures



Finite impulse Response (FIR)

No feedback

Order -4



Infinite impulse Response (IIR)

Feedback

Order - 2

SU2010-CS

Algo2Arch

25

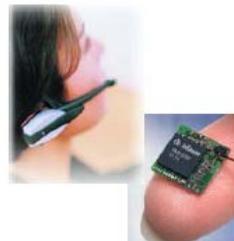
Different applns different demands



Flexibility
Complexity
Processors
FPGAs



Low power
Low cost
Flexibility
Processors
ASICs



Lower power
Lower cost
ASICs
Processors

SU2010-CS

Algo2Arch

26

Standard processors or special purpose ?

```

    graph TD
        Algorithm((Algorithm)) --> StandardProcessor([Standard Processor])
        Algorithm --> SpecialPurpose([Special Purpose])
        StandardProcessor <--> SpecialPurpose
    
```

Standard Processor

- Programmable/Flexible
- Short design time/TTM
- Low price?

Special Purpose

- High calculation capacity
- Low power consumption
- Low price at volume
– What is volume?

SU2010-CS Algo2Arch 27

Architectural options

OTS (Off The Shelf) processors

- Programmable microprocessors or DSP
- Based on generic computational units, for DSPs usually MAC
- Prefabbed or IP cores

Time-multiplexed application specific processors

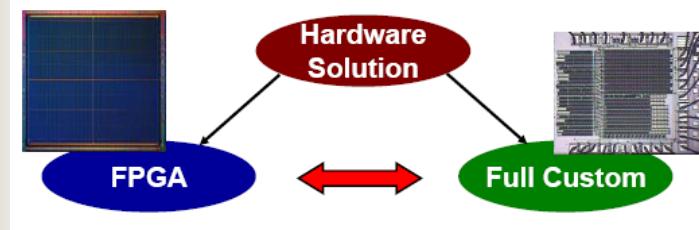
- Several algorithmic operations performed on same hardware unit
- Trades reduced HW for longer computation time

Hardware mapped architectures

- One (or more) hardware unit per algorithmic operation
- High hardware cost and high throughput

SU2010-CS Algo2Arch 28

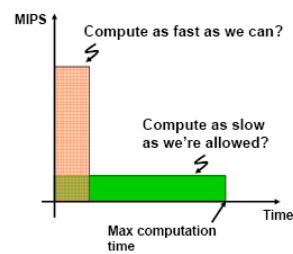
Hardware options



FPGA - Field Programmable gate arrays

ASIC – Application Specific Integrated Circuit ASIC

Key design issue today - energy



Utilising the computational time
clock frequency
supply voltage
Sleep modes



Algorithm 2 Architecture

- Which structure gives optimal performance, energy and area?
- How to get from a signal processing algorithm to an **EFFICIENT** implementation using
 - Different numbering systems
 - Pipelining
 - Parallelism
 - Retiming, Scheduling
 - Strength reduction, i.e. complexity of operations.
 - etc, etc,...
- in a structured way!

SU2010-CS

Algo2Arch

31



An example

- Datapath and control path

SU2010-CS

Algo2Arch

32



DSP basic operations



FIR $y_n = \sum_{k=0}^{N-1} h_k \cdot x_{n-k}$

IIR $y_n = \sum_{k=0}^M a_k \cdot x_{n-k} + \sum_{k=1}^N b_k \cdot y_{n-k}$

Transforms
FFT,DCT $X_m = \sum_{n=0}^{N-1} x_n \cdot e^{-j2\pi m n / N}$

Decomposition (SVD, LU, QR) - Matrix operations

Arithmetic operations ? x, +, shifts

Non terminating – data flow , sampling rate

SU2010-CS *Algo2Arch* 33



Number representation



- Binary number system
 - Integers eg. 0110 → 6 (decimal)
 - Two's complement 1010 → -6 (decimal) (4 bit)
 - Fractional values
 - Divide by 2^{n-1} → range [-1, +1]
 - n bits : $[2^{n-1}, 2^{n-1}-1]$
 - 8 bits : ?
 - Fractional : -1, +127/128
 - Resolution : 1/256 of range :[-1, +1]

SU2010-CS *Algo2Arch* 34

Floating point numbers

- 32 bits

S(1)	Exponent (8)	Mantissa (23)
------	--------------	---------------

- Number magnitude: $(1\text{-mantissa}) \times 2^{\text{exponent}-127}$
 - Smallest +ve: $(1.00\ldots0) \times 2^{-126}$
 - Largest +ve: $(1.11\ldots1) \times 2^{127} \sim 2^{128}$
 - Mantissa – determines resolution
 - Exponent – determines dynamic range
- Much wider dynamic range at the cost of energy, area, computation time
- Energy efficient implementation fixed point used

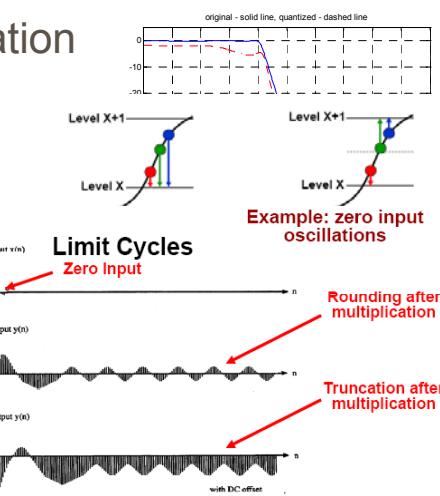
SU2010-CS

Algo2Arch

35

Finite word length effects

- Coefficient quantisation
- Signal quantisation
 - Round off noise
 - Limit cycles
- Scaling
 - Adjust signal to hardware
- Saturation arithmetic



SU2010-CS

Algo2Arch

36



Implementation Flow

- High level description - Specifications
 - Equations
 - Matlab, C/C++
- Architecture
 - Building blocks, speed and area
 - Block diagram
 - Optimisations: #bits, # coefficients
 - Hardware Description languages – Verilog, VHDL
- Compile (Synthesize)
 - Mapping to gates
 - Place & Route

SU2010-CS

Algo2Arch

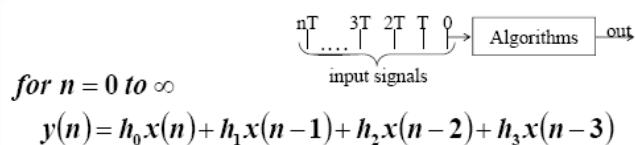
37



Implementation of FIR filters

$$y_n = \sum_{k=0}^3 h_k \cdot x_{n-k}$$

Algorithm non-terminating



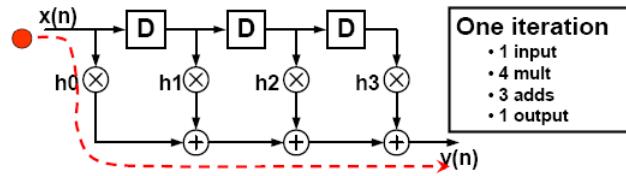
One execution of the loop – one iteration
Iteration period = time for one iteration

SU2010-CS

Algo2Arch

38

Direct form 4 tap Fir filter



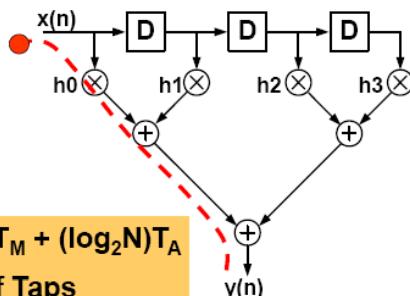
Critical path = longest path without delay element

Critical path sets bound on clock frequency

$$T_{clk} \geq T_{critical}$$

$$T_{critical} = T_M + (N-1)T_A$$

FIR with adder tree



$$T_{critical} = T_M + (\log_2 N)T_A$$

N = Nr. of Taps

Critical path



- Critical path
 - between delays
 - input to delay
 - delay to output
 - input to output
- Sampling rate = no of samples processed/sec
- Latency = time difference between output and corresponding input
 - Combinatorial logic = gate delays
 - Sequential logic = no of clock cycles

SU2010-CS

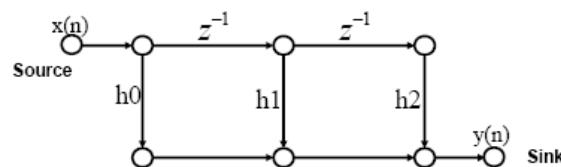
Algo2Arch

41

Signal Flow Graph



- Nodes: represent computations and/or task
- Directed edge (j, k) : denotes a linear transformation from the input signal at node j to the output signal at node k
 - in digital usually limited to delays and constant multipliers
- Source = no entering edge
- Sink = only entering edges

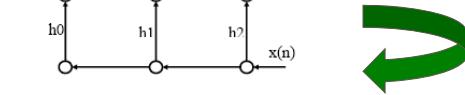
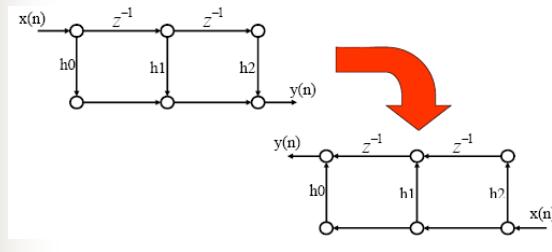


SU2010-CS

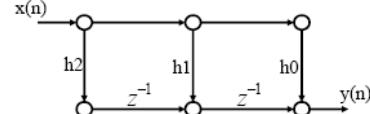
Algo2Arch

42

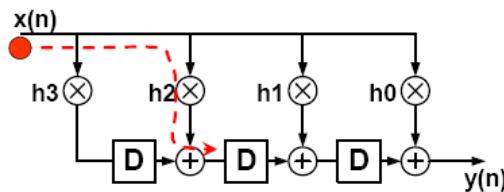
Transposition



Reverse the direction of all edges
Exchange input and output



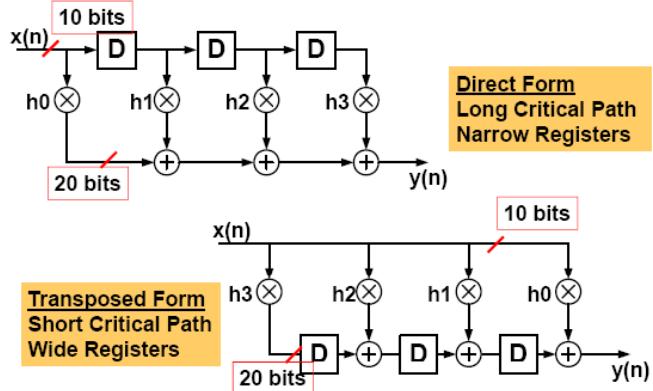
Transposed FIR filter



$$T_{\text{Critical}} = T_M + T_A$$

$$N = \text{Nr. of Taps}$$

Fixed point implementation

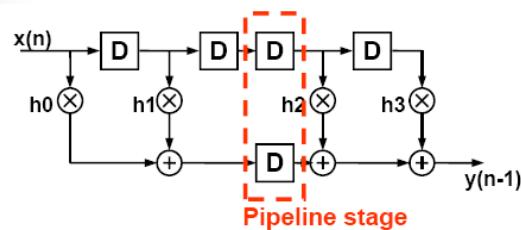


SU2010-CS

Algo2Arch

45

Pipelining



Pipelining increases latency

More registers can bring down Critical Path delay to $\max(T_M, T_A)$

SU2010-CS

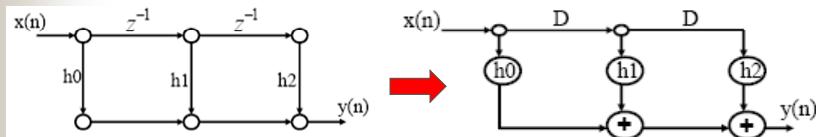
Algo2Arch

46

Data Flow Graph



- DFGs capture the data-driven property of DSP algorithm
- Data-Flow Graph – nodes represent computations (or functions)
 - directed edges represent data flow with non-negative number of delays
 - a node can execute whenever all its input data are available.
 - Each edge describes a precedence constraint between two nodes in DFG:
- Intra-iteration precedence constraint: if the edge has zero delays
- Inter-iteration precedence constraint: if the edge has one or more delays

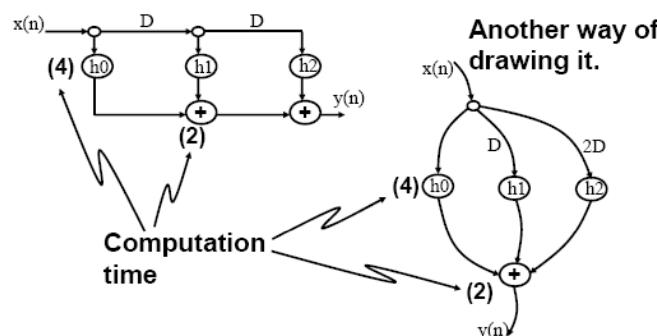


SU2010-CS

Algo2Arch

47

DFG



SU2010-CS

Algo2Arch

48

Dependence Graph



- Like DFG a combination of nodes (tasks) with directed edges which indicate precedence constraints
- Unlike DFG nodes are not reused. Instead in each iteration a new node is created
- DG has an explicit time axis, hence delays are not indicated
- Extremely modular regular arrays

SU2010-CS

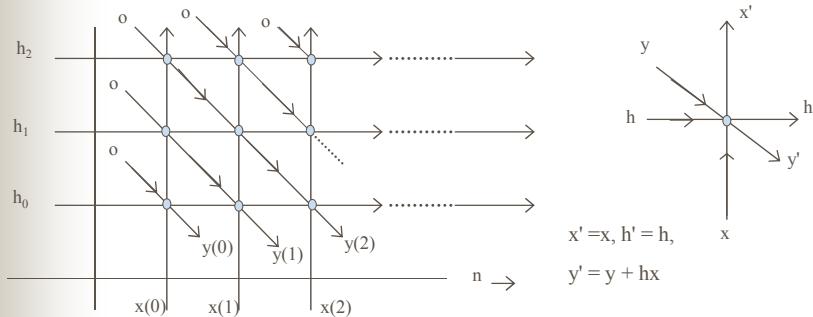
Algo2Arch

49

FIR Filter - DG



$$y(n) = h_0 x(n) + h_1 x(n-1) + h_2 x(n-2)$$



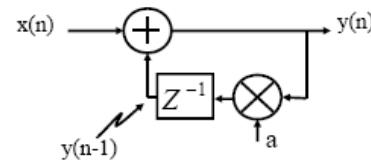
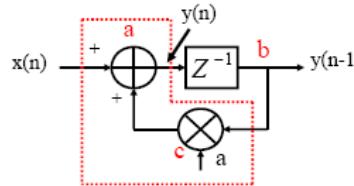
SU2010-CS

Algo2Arch

50

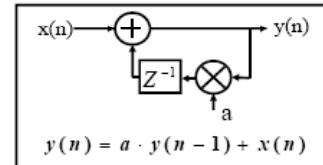
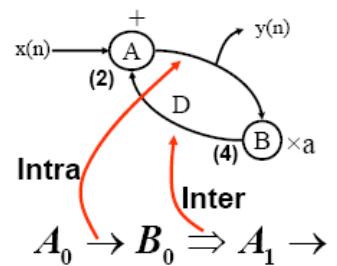
Recurrence equations

$$y(n) = a \cdot y(n-1) + x(n)$$

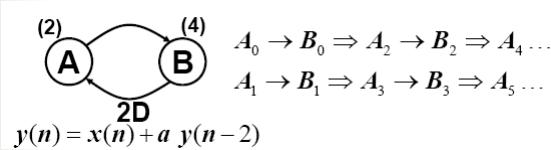


Time for one complete cycle = $T_M + T_A$ – Iteration period
 Sampling time $T_S \geq T_M + T_A$

Precedence constraints



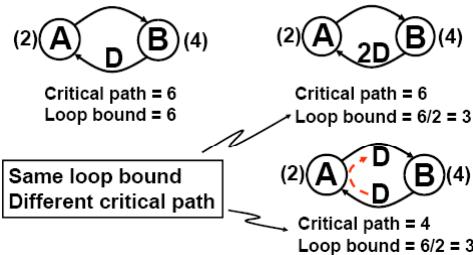
$$y(n) = a \cdot y(n-1) + x(n)$$



$$y(n) = x(n) + a \cdot y(n-2)$$

Loop Bound

$$\text{Loop bound} = \frac{\text{Total execution time on loop}}{\text{Total delays in loop}}$$



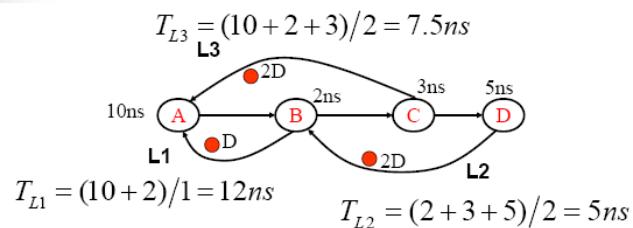
SU2010-CS

Algo2Arch

53

Iteration Bound

Iteration Bound = Maximum loop bound



Delay free loops cannot be computed

SU2010-CS

Algo2Arch

54



Iteration bound

- Clock period is lower bounded by the critical path computation time
- Sampling period is lower bounded by the iteration bound regardless of computational resources available
- Iteration bound
 - Depends on technology
 - Derivation of graph from equations

SU2010-CS

Algo2Arch

55



Pipelining and Parallelism

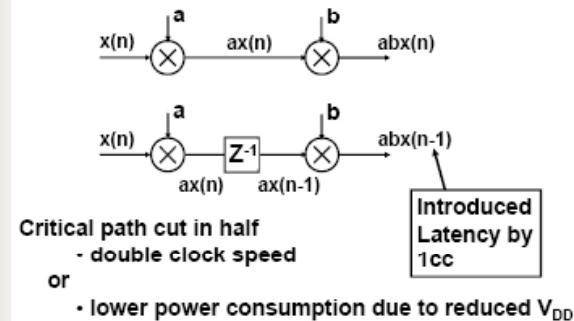
- Pipelining
 - Splits the logic path by introducing pipeline registers
 - leads to a reduction of the critical path but introduce latency
 - Either increases the clock speed (or sampling speed) or reduces the power consumption at same speed in a DSP system
- Parallel Processing
 - Multiple outputs are computed in parallel in a clock period
 - The effective sampling speed is increased by the level of parallelism
 - Can also be used to reduce the power consumption

SU2010-CS

Algo2Arch

56

Pipelining

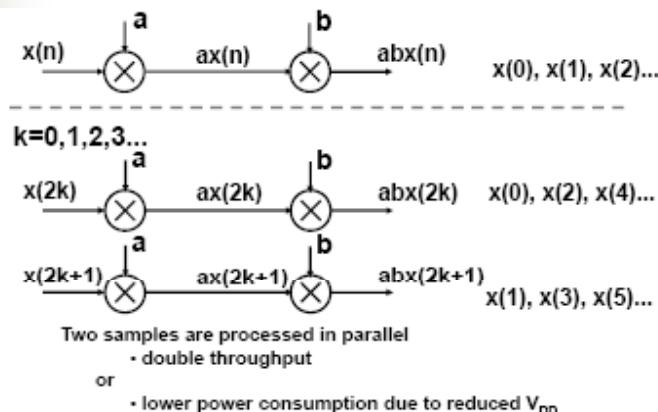


SU2010-CS

Algo2Arch

57

Parallel processing



SU2010-CS

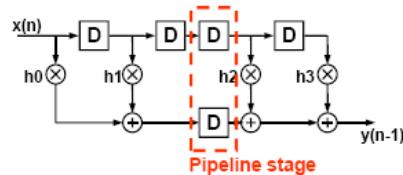
Algo2Arch

58

Pipelining



- Cutset - A set of edges that if removed, or cut, results in two disjoint graphs.
- Feedforward Cutset - if data is moved in forward direction on all cutsets
- Pipelining – Adding delays at feedforward cutsets
- Reduces critical path delay
- Introduces latency, no change in functionality

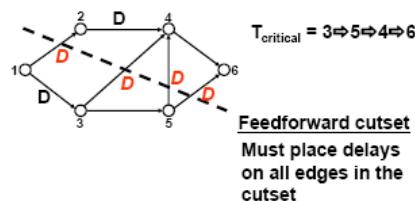


SU2010-CS

Algo2Arch

59

Pipelining



$$T_{node} = 1$$

$$CP_{old} = 4$$

$$CP_{new} = 2$$

SU2010-CS

Algo2Arch

60

Parallel FIR filter

SISO FIR $y(n) = b_0x(n) + b_1x(n-1) + b_2x(n-2)$

MIMO FIR $\begin{cases} y(2k) = b_0x(2k) + b_1x(2k-1) + b_2x(2k-2) \\ y(2k+1) = b_0x(2k+1) + b_1x(2k) + b_2x(2k-1) \end{cases}$

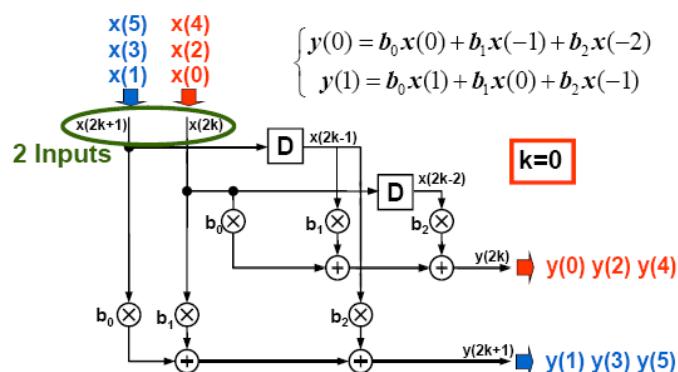


SU2010-CS

Algo2Arch

61

Parallel 3 tap FIR



SU2010-CS

Algo2Arch

62

Pipelining & Parallelism



- Critical path unchanged in parallel system
 - $T_{clock} \geq T_M + 2T_A$
 - $T_{iter} = T_s = T_{clock}/L$
- Parallel system $T_{clock} \neq T_s$
- Pipeline system $T_{clock} = T_s$
- By combining parallel processing (block size: L) and pipelining (pipelining stage: M), the sample period can be reduced to
- $T_{iter} = T_s = T_{clock}/LM$

SU2010-CS

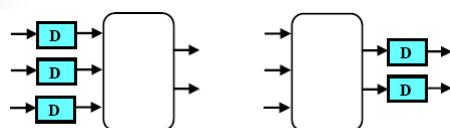
Algo2Arch

63

Retiming



Moving delays in the system



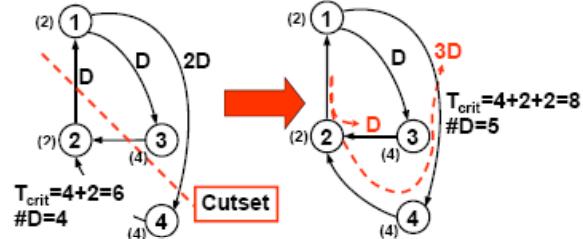
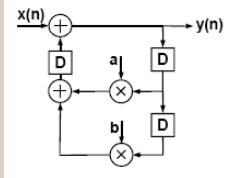
- Modifies Critical path delay, no of registers
- Retiming does not change
 - delay in a loop
 - iteration bound

SU2010-CS

Algo2Arch

64

Cutset retiming



Cutset Retiming
Add delays to edges
going one way and
remove from ones going
the other.

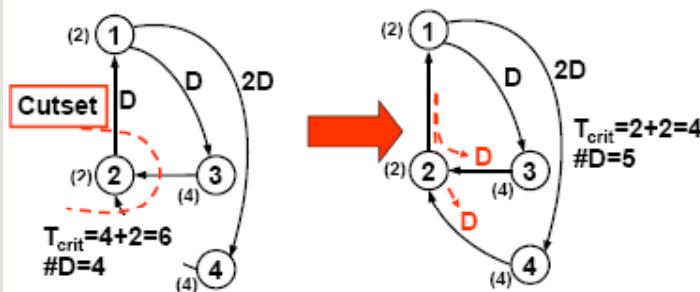
SU2010-CS

Algo2Arch

65

Node retiming

Cut set around a node



SU2010-CS

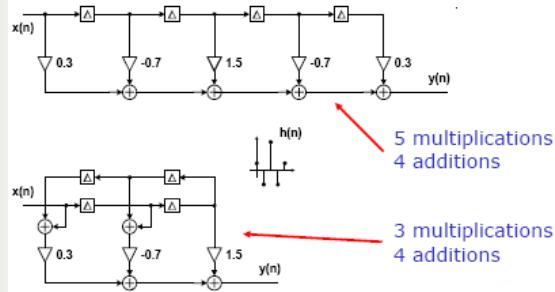
Algo2Arch

66

Algorithmic Strength reduction



- Reduces no of strong operations



SU2010-CS

Algo2Arch

67

Fixed coefficient multiplication



Conventional Multiplication

$$\begin{array}{r}
 Z = X \cdot Y \\
 \hline
 \begin{array}{ccccccccc}
 & X_3 & X_2 & X_1 & X_0 \\
 & Y_3 & Y_2 & Y_1 & Y_0 \\
 \hline
 X_3 \cdot Y_0 & X_2 \cdot Y_0 & X_1 \cdot Y_0 & X_0 \cdot Y_0 \\
 X_3 \cdot Y_1 & X_2 \cdot Y_1 & X_1 \cdot Y_1 & X_0 \cdot Y_1 \\
 X_3 \cdot Y_2 & X_2 \cdot Y_2 & X_1 \cdot Y_2 & X_0 \cdot Y_2 \\
 X_3 \cdot Y_3 & X_2 \cdot Y_3 & X_1 \cdot Y_3 & X_0 \cdot Y_3 \\
 \hline
 Z_7 & Z_6 & Z_5 & Z_4 & Z_3 & Z_2 & Z_1 & Z_0
 \end{array}
 \end{array}$$

Constant multiplication (become hardwired shifts and adds)

$$\begin{array}{r}
 Z = X \cdot (1001)_2 \\
 \hline
 \begin{array}{ccccccccc}
 & X_3 & X_2 & X_1 & X_0 \\
 & 1 & 0 & 0 & 1 \\
 \hline
 X_3 & X_2 & X_1 & X_0 \\
 X_3 & X_2 & X_1 & X_0 \\
 \hline
 Z_7 & Z_6 & Z_5 & Z_4 & Z_3 & Z_2 & Z_1 & Z_0
 \end{array}
 \end{array}$$

$$\begin{array}{l}
 Y = (1001)_2 = 2^3 + 2^0 \\
 \hline
 X \xrightarrow{\ll 3} \text{shifted } X \xrightarrow{+Y} Z
 \end{array}$$

shifts using wiring

SU2010-CS

Algo2Arch

68



Complex multiplication

- $(a+jb)(x+jy)$
- Can you do it with 3 multipliers and few extra adders?

SU2010-CS

Algo2Arch

69



Architectural synthesis

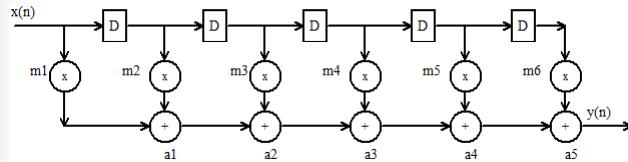
- Given a DFG, resources
 - Allocation
 - Allocate enough resources to solve the problem
 - Binding
 - Which operation happens on which resource
 - Scheduling
 - When should each operation take place

SU2010-CS

Algo2Arch

70

6 tap FIR Filter

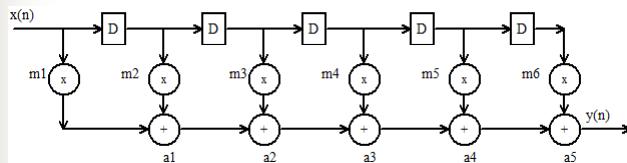


Cycle	Operations	Resources
1	m1, m2,m3,m4,m5,m6	6 multi
2	a1	1 adder
3	a2	1 adder
4	a3	1 adder`
5	a4	1 adder
6	a5	1 adder

Resource Utilisation

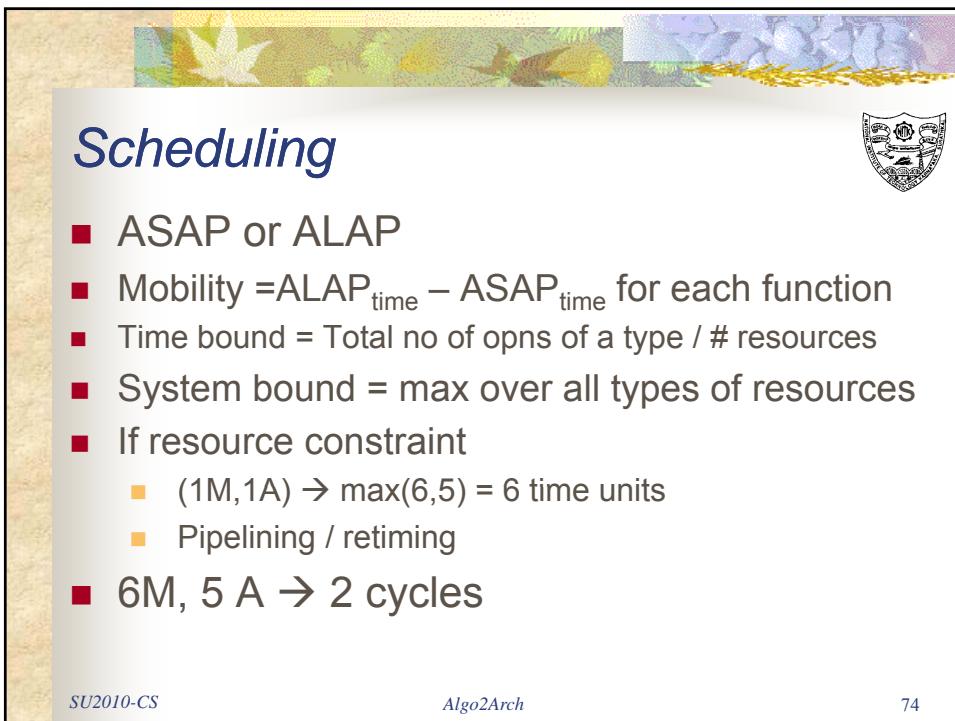
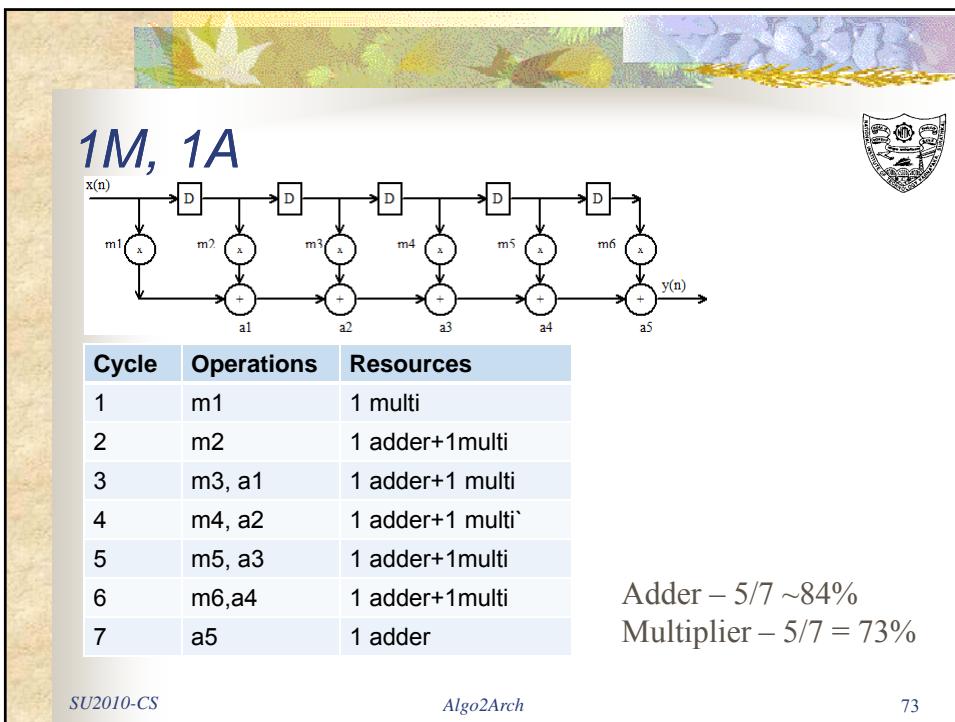
Adder – 5/6 ~83%
 Multiplier – 6/36 ~ 16.67%
 m1, m2 have to be in cycle 1

Alternate – 2 M +1 A

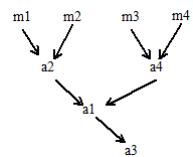
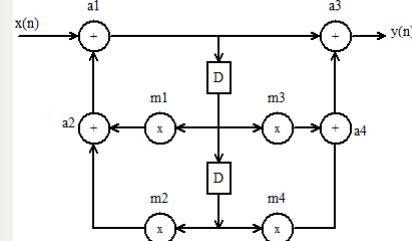


Cycle	Operations	Resources
1	m1, m2	2 multi
2	a1,m3,m4	1 adder+1multi
3	a2,m5,m6	1 adder+1 multi
4	a3	1 adder`
5	a4	1 adder
6	a5	1 adder

Adder – 5/6 ~83%
 Multiplier – 6/12 = 50%



IIR filter

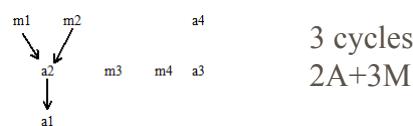
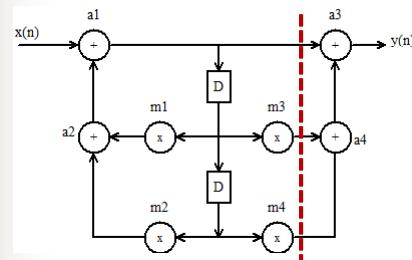


SU2010-CS

Algo2Arch

75

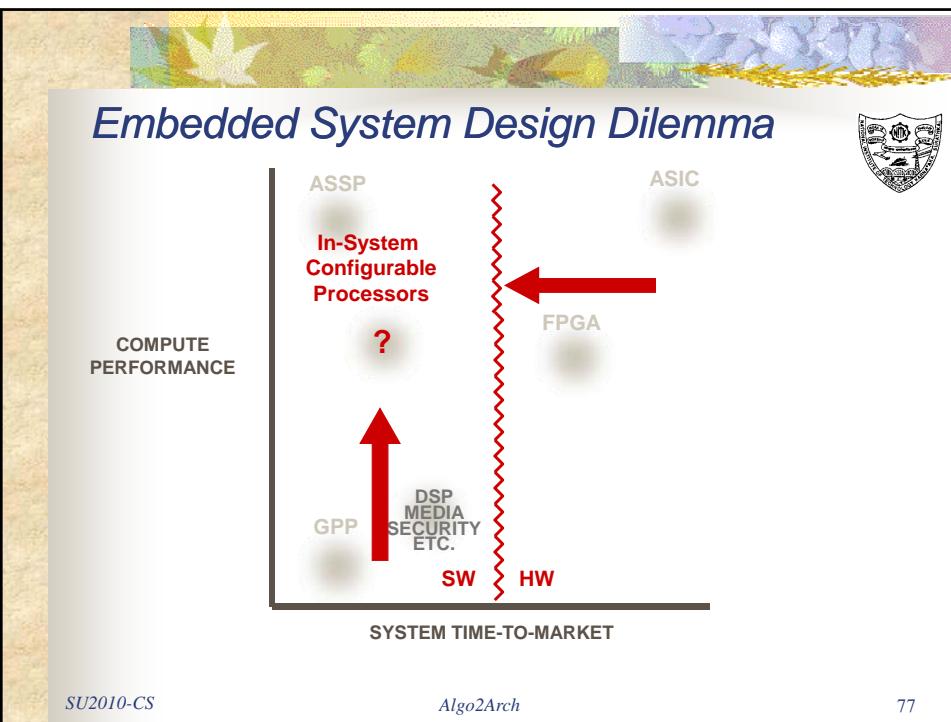
IIR filter



SU2010-CS

Algo2Arch

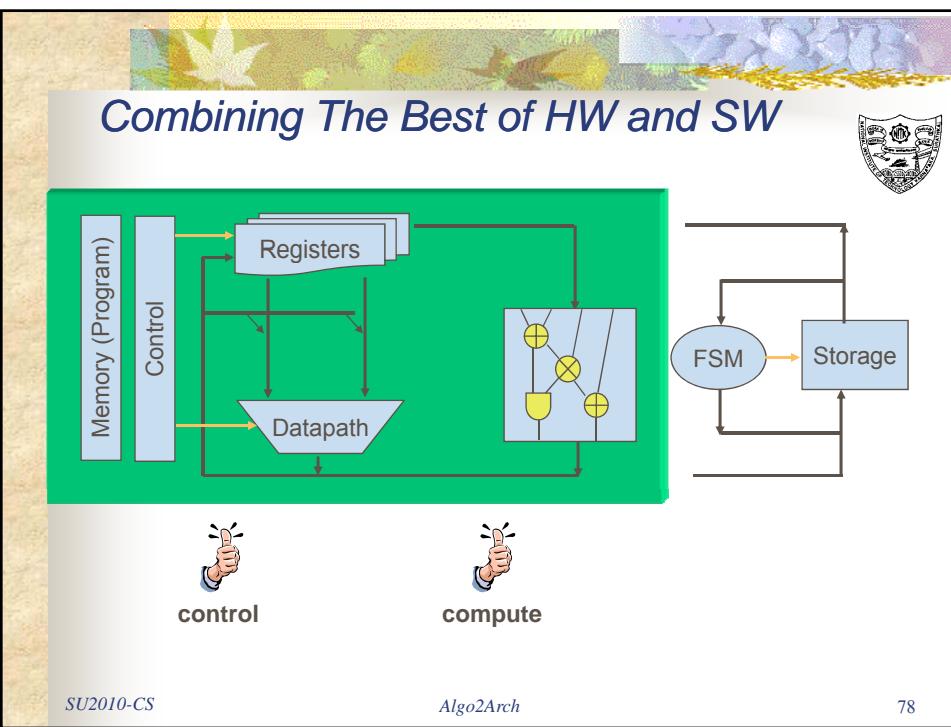
76



SU2010-CS

Algo2Arch

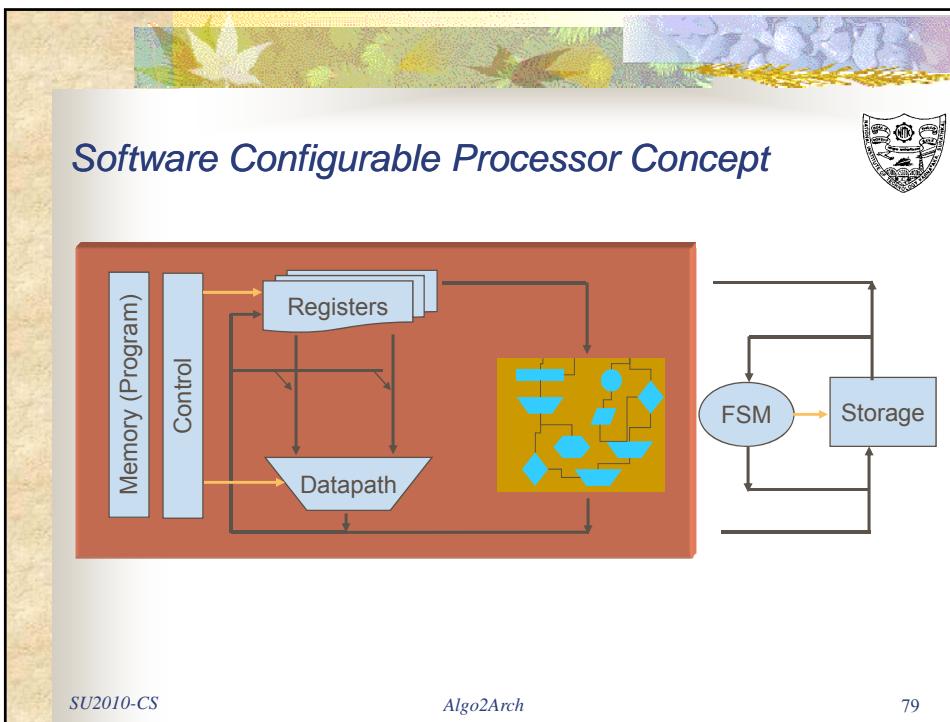
77



SU2010-CS

Algo2Arch

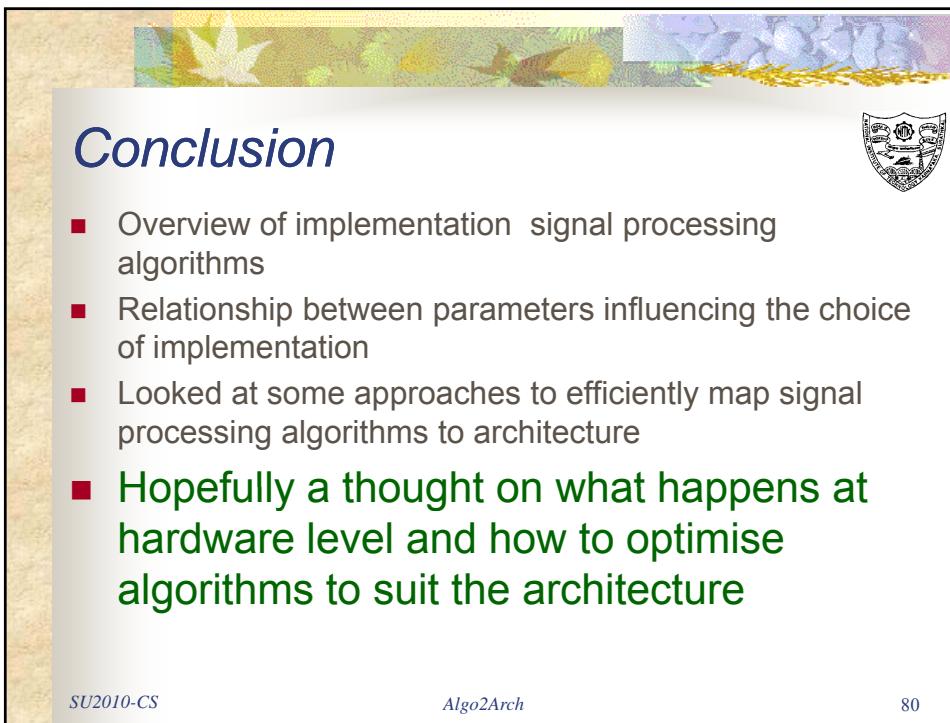
78



SU2010-CS

Algo2Arch

79



SU2010-CS

Algo2Arch

80



References



- K.K. Parhi, *VLSI Digital Signal Processing Systems*, Wiley 1999

SU2010-CS

Algo2Arch

81

