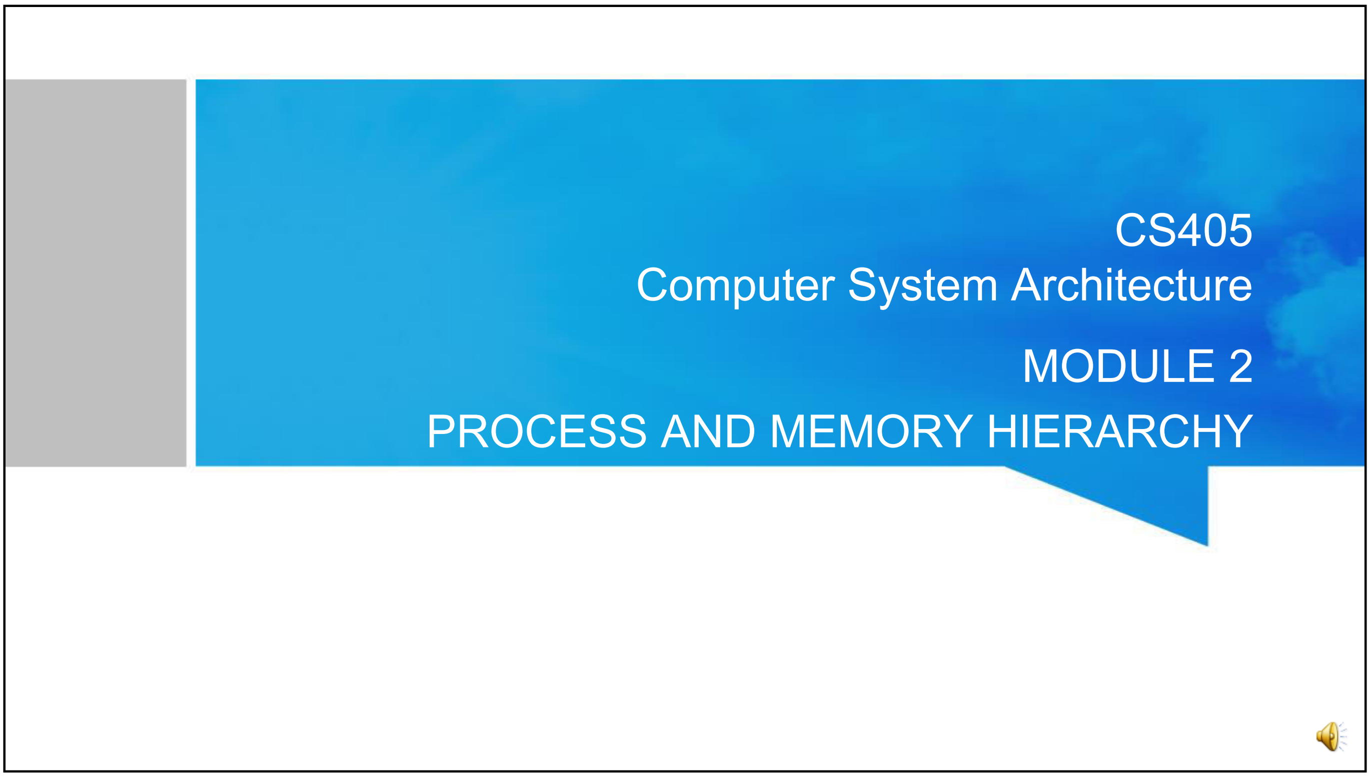

CS405
Computer System Architecture

MODULE 2
PROCESS AND MEMORY HIERARCHY

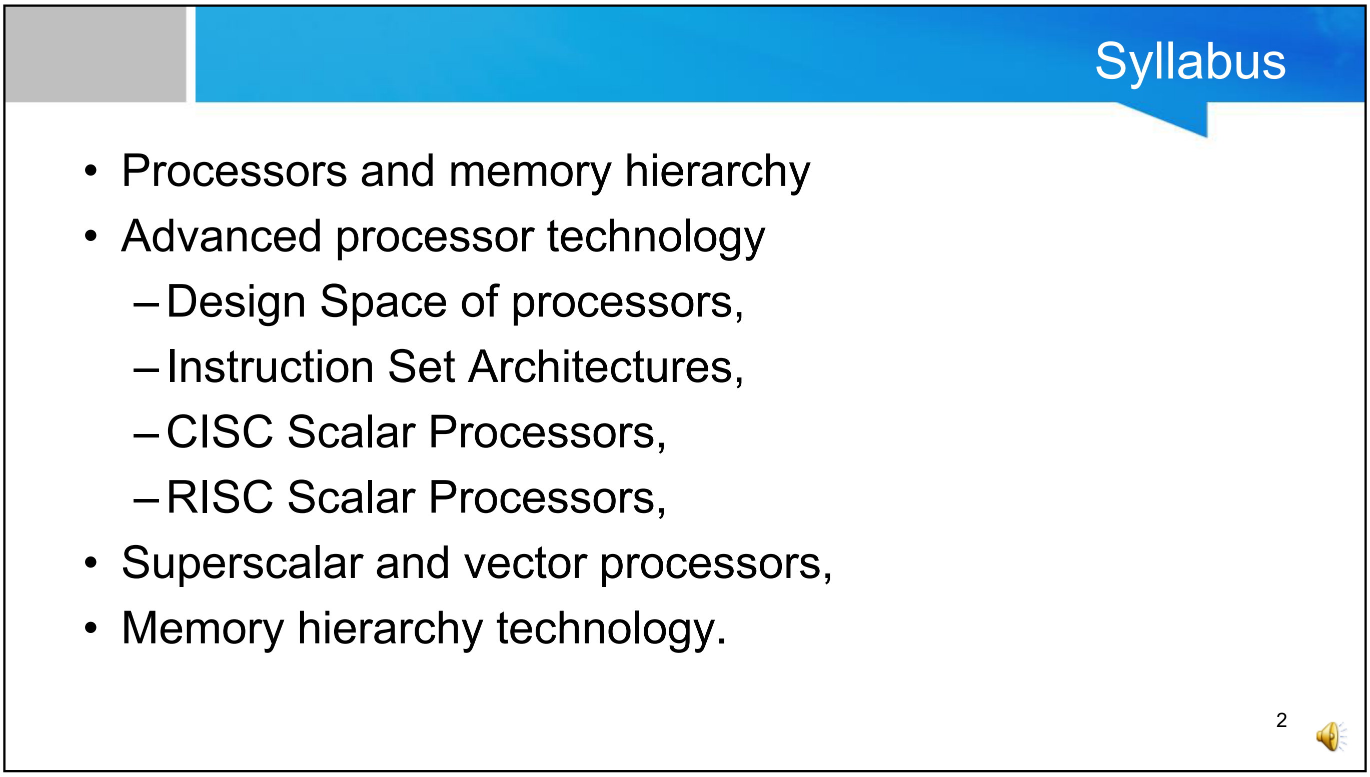

Syllabus

- Processors and memory hierarchy
- Advanced processor technology
  - Design Space of processors,
  - Instruction Set Architectures,
  - CISC Scalar Processors,
  - RISC Scalar Processors,
- Superscalar and vector processors,
- Memory hierarchy technology.

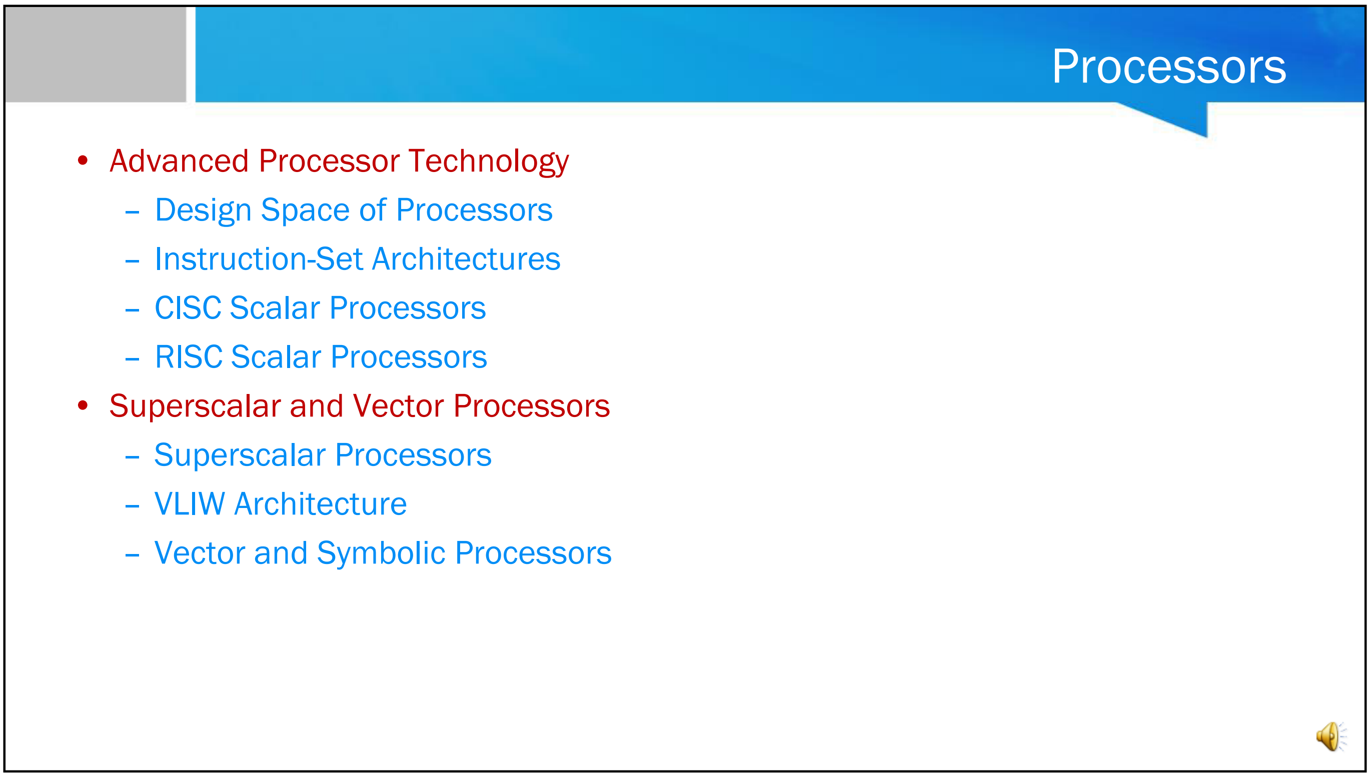

Processors

- Advanced Processor Technology
  - Design Space of Processors
  - Instruction-Set Architectures
  - CISC Scalar Processors
  - RISC Scalar Processors
- Superscalar and Vector Processors
  - Superscalar Processors
  - VLIW Architecture
  - Vector and Symbolic Processors

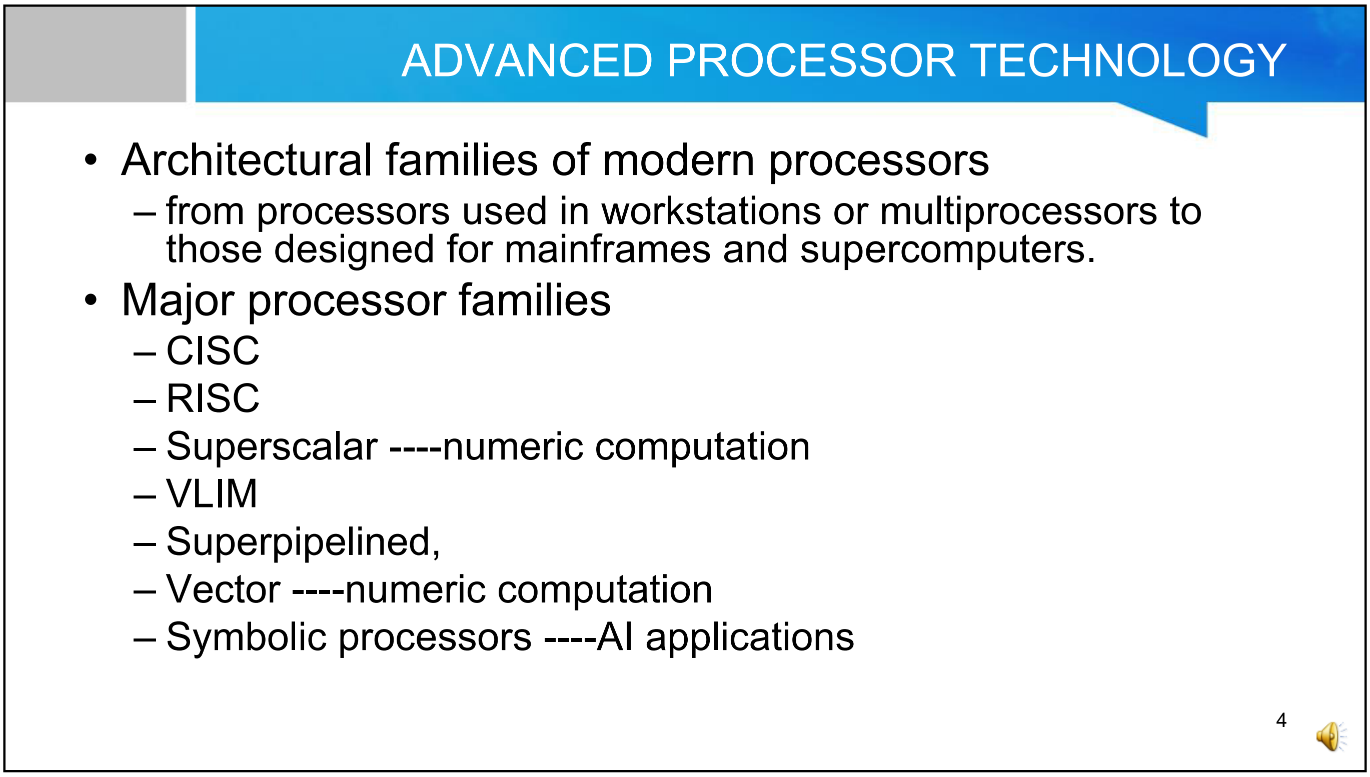

ADVANCED PROCESSOR TECHNOLOGY

- Architectural families of modern processors
  - from processors used in workstations or multiprocessors to those designed for mainframes and supercomputers.
- Major processor families
  - CISC
  - RISC
  - Superscalar ----numeric computation
  - VLIM
  - Superpipelined,
  - Vector ----numeric computation
  - Symbolic processors ----AI applications

Design Space of Processors

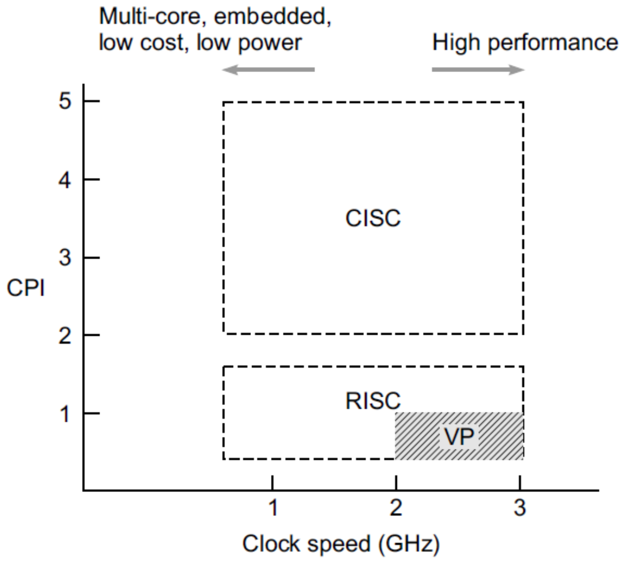

Fig. 4.1 CPI versus processor clock speed of major categories of processors

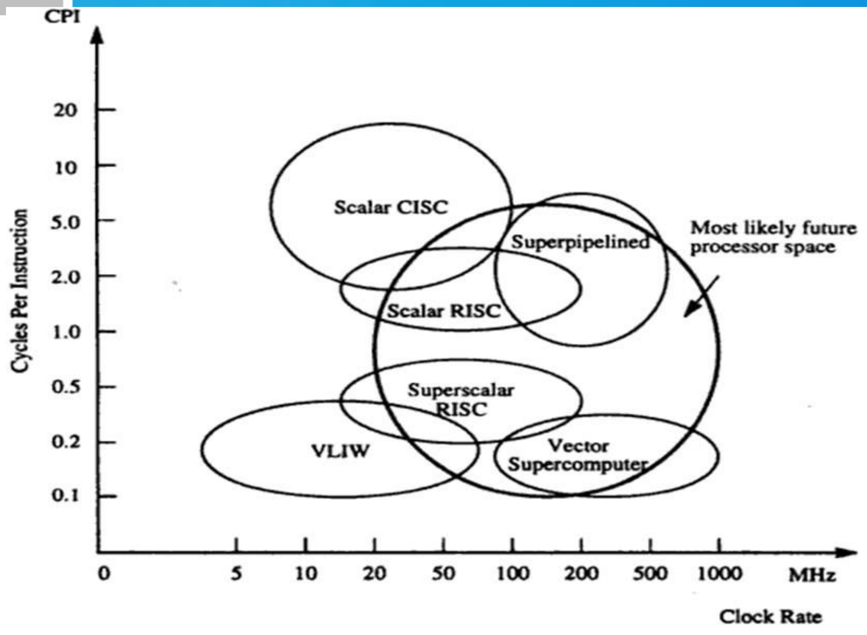

Design space of modern processor families.

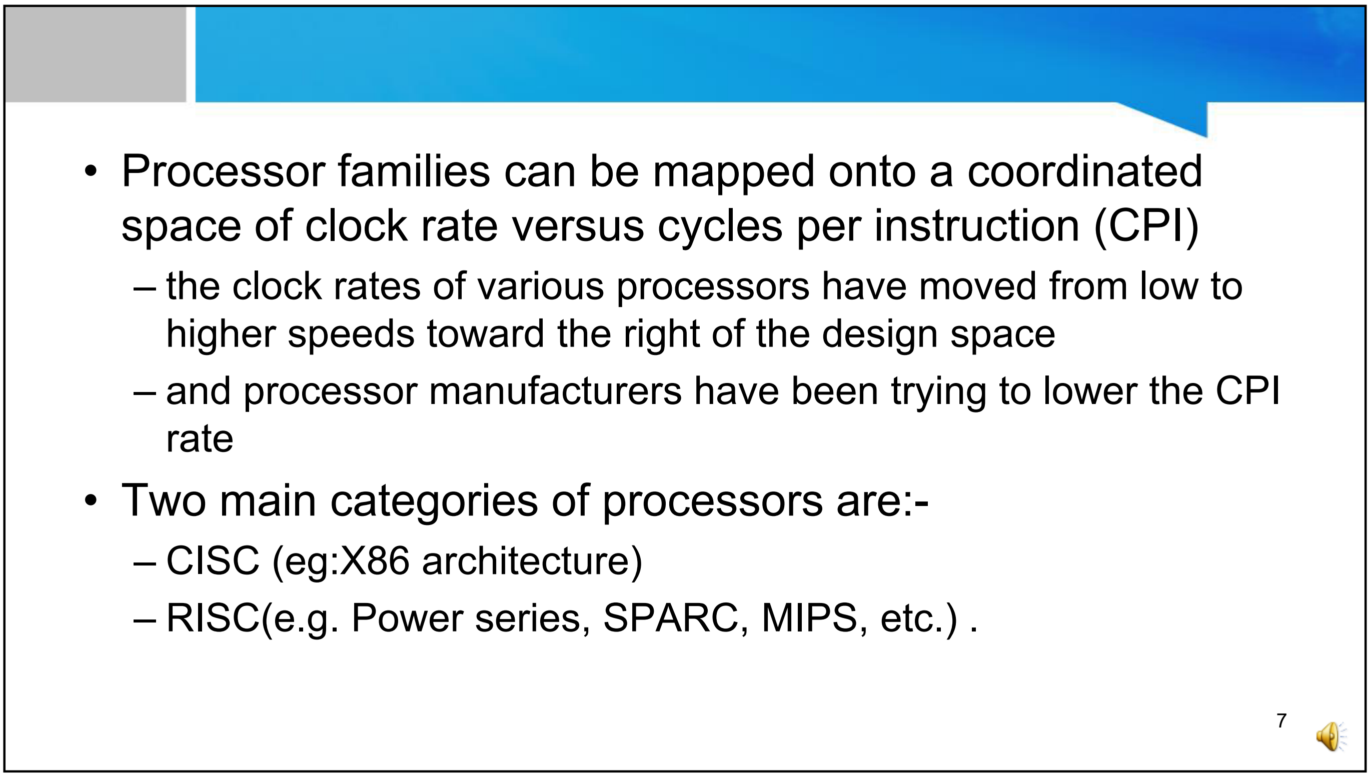

- Processor families can be mapped onto a coordinated space of clock rate versus cycles per instruction (CPI)
  - the clock rates of various processors have moved from low to higher speeds toward the right of the design space
  - and processor manufacturers have been trying to lower the CPI rate
- Two main categories of processors are:-
  - CISC (eg:X86 architecture)
  - RISC(e.g. Power series, SPARC, MIPS, etc.) .

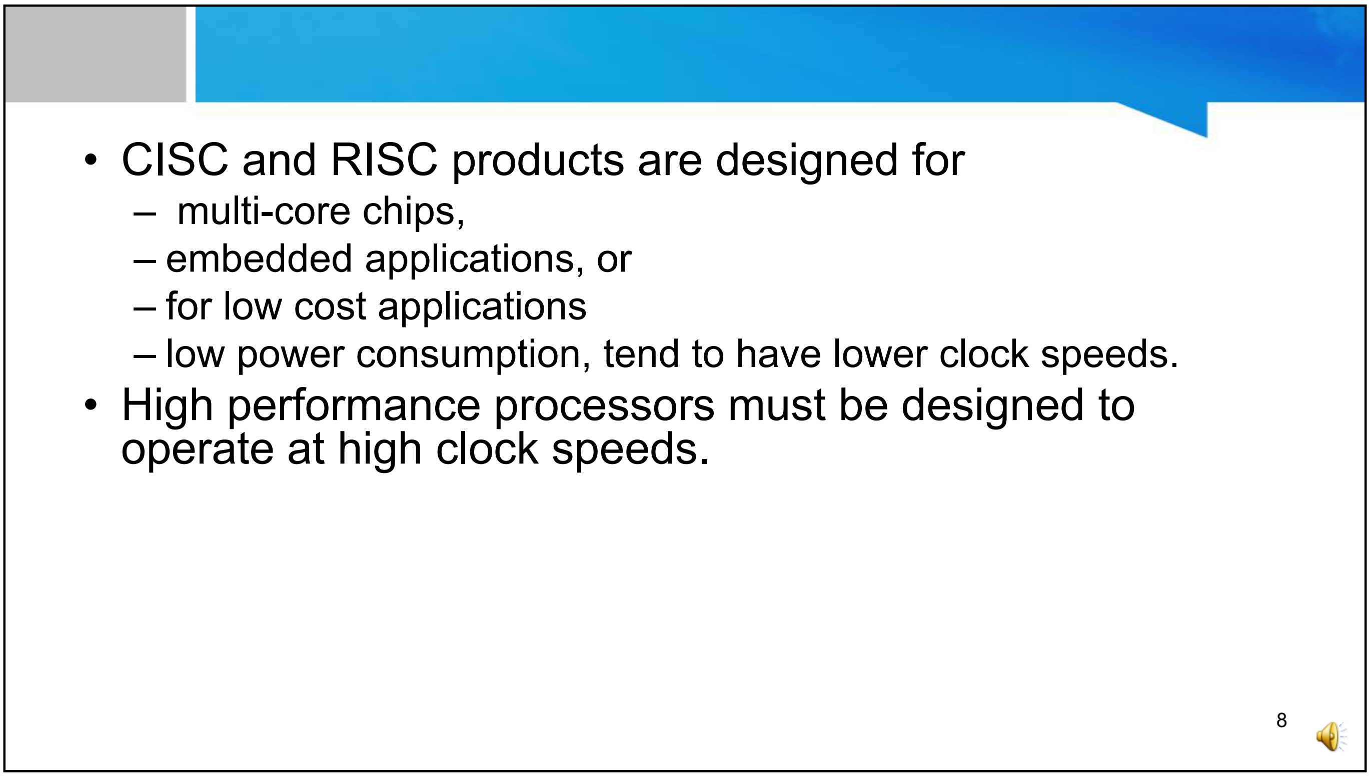

- CISC and RISC products are designed for
  - multi-core chips,
  - embedded applications, or
  - for low cost applications
  - low power consumption, tend to have lower clock speeds.
- High performance processors must be designed to operate at high clock speeds.

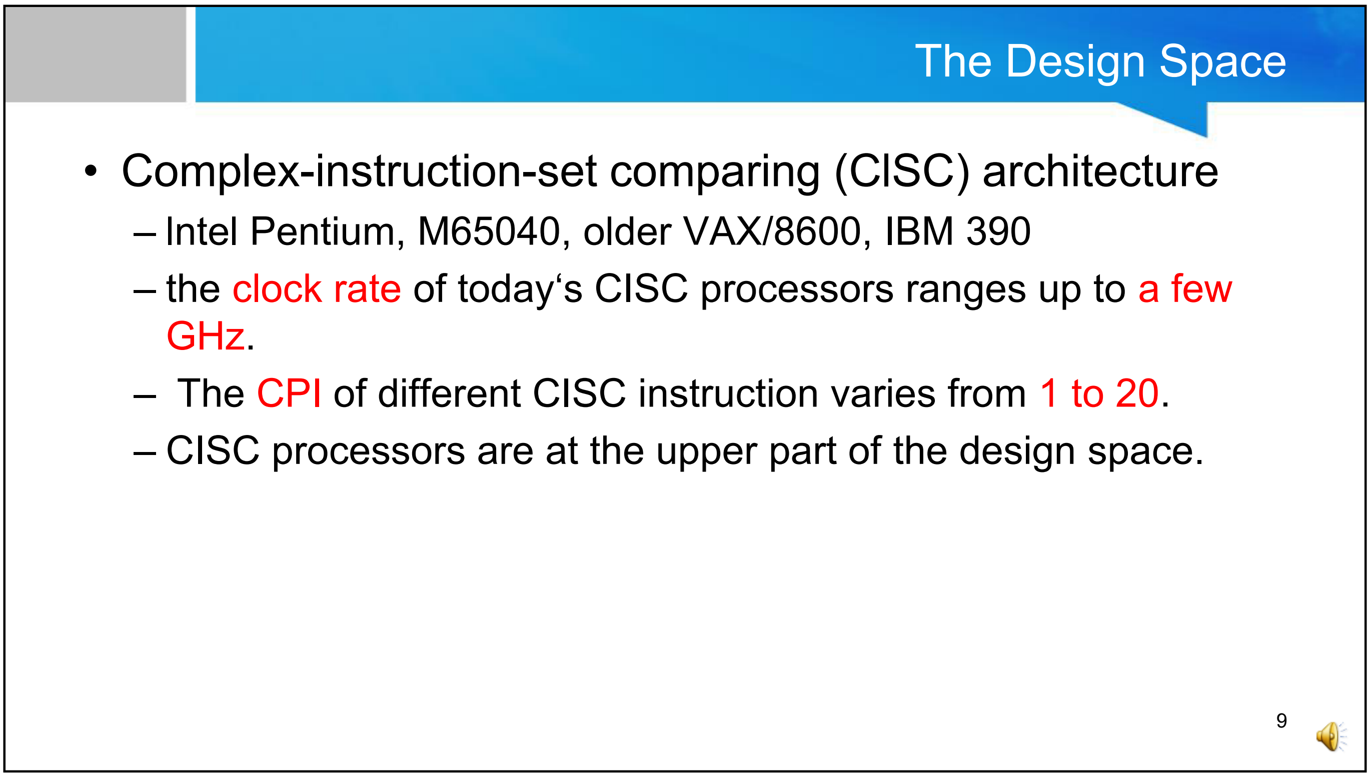

The Design Space

- Complex-instruction-set comparing (CISC) architecture
  - Intel Pentium, M65040, older VAX/8600, IBM 390
  - the clock rate of today's CISC processors ranges up to a few GHz.
  - The CPI of different CISC instruction varies from 1 to 20.
  - CISC processors are at the upper part of the design space.

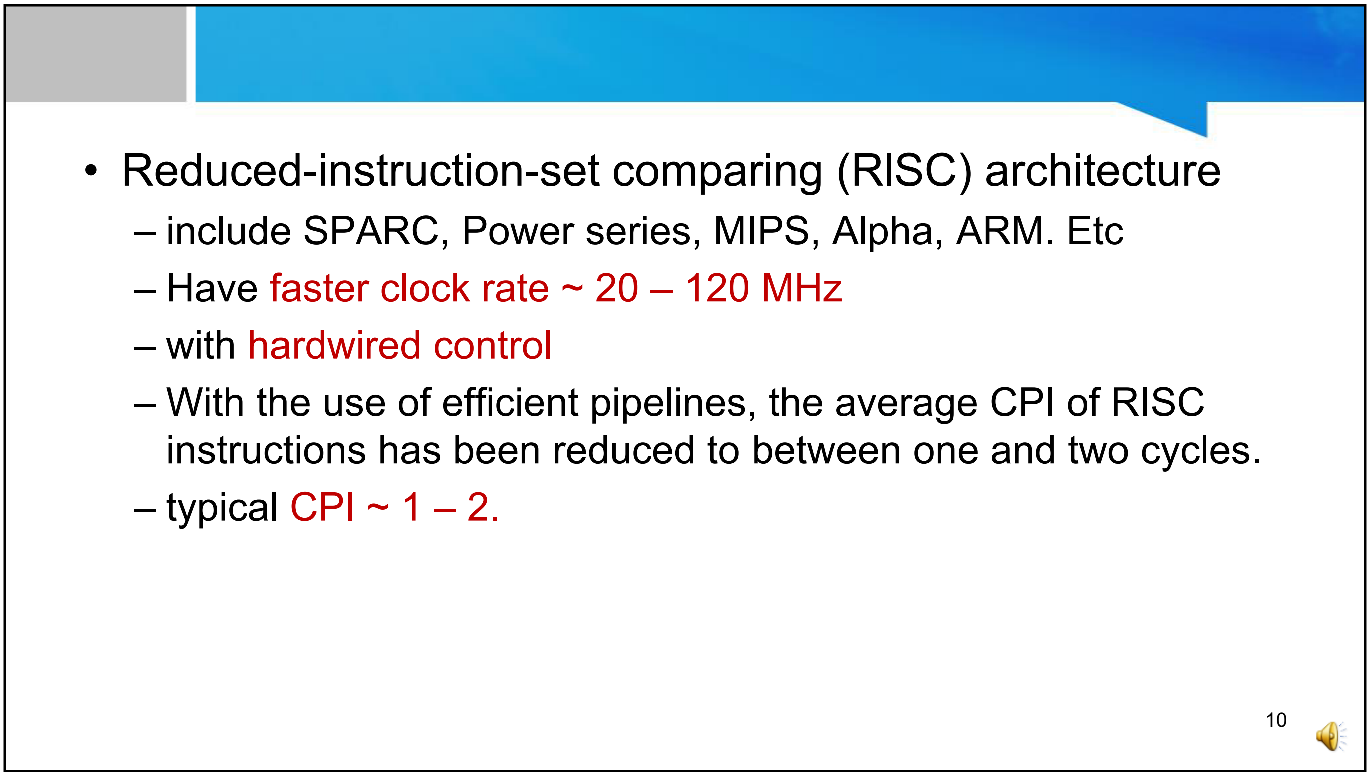

- Reduced-instruction-set comparing (RISC) architecture
  - include SPARC, Power series, MIPS, Alpha, ARM. Etc
  - Have faster clock rate ~ 20 – 120 MHz
  - with hardwired control
  - With the use of efficient pipelines, the average CPI of RISC instructions has been reduced to between one and two cycles.
  - typical CPI ~ 1 – 2.

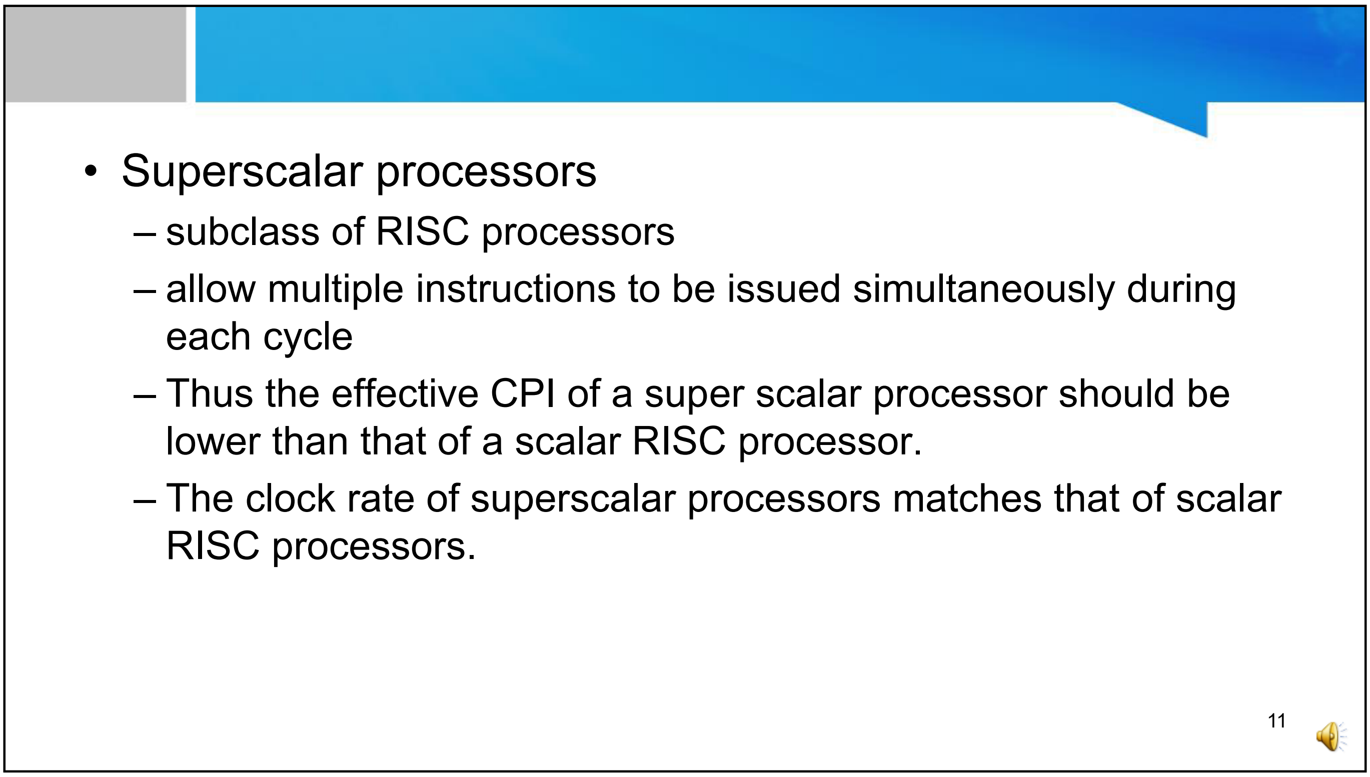

- Superscalar processors
  - subclass of RISC processors
  - allow multiple instructions to be issued simultaneously during each cycle
  - Thus the effective CPI of a super scalar processor should be lower than that of a scalar RISC processor.
  - The clock rate of superscalar processors matches that of scalar RISC processors.

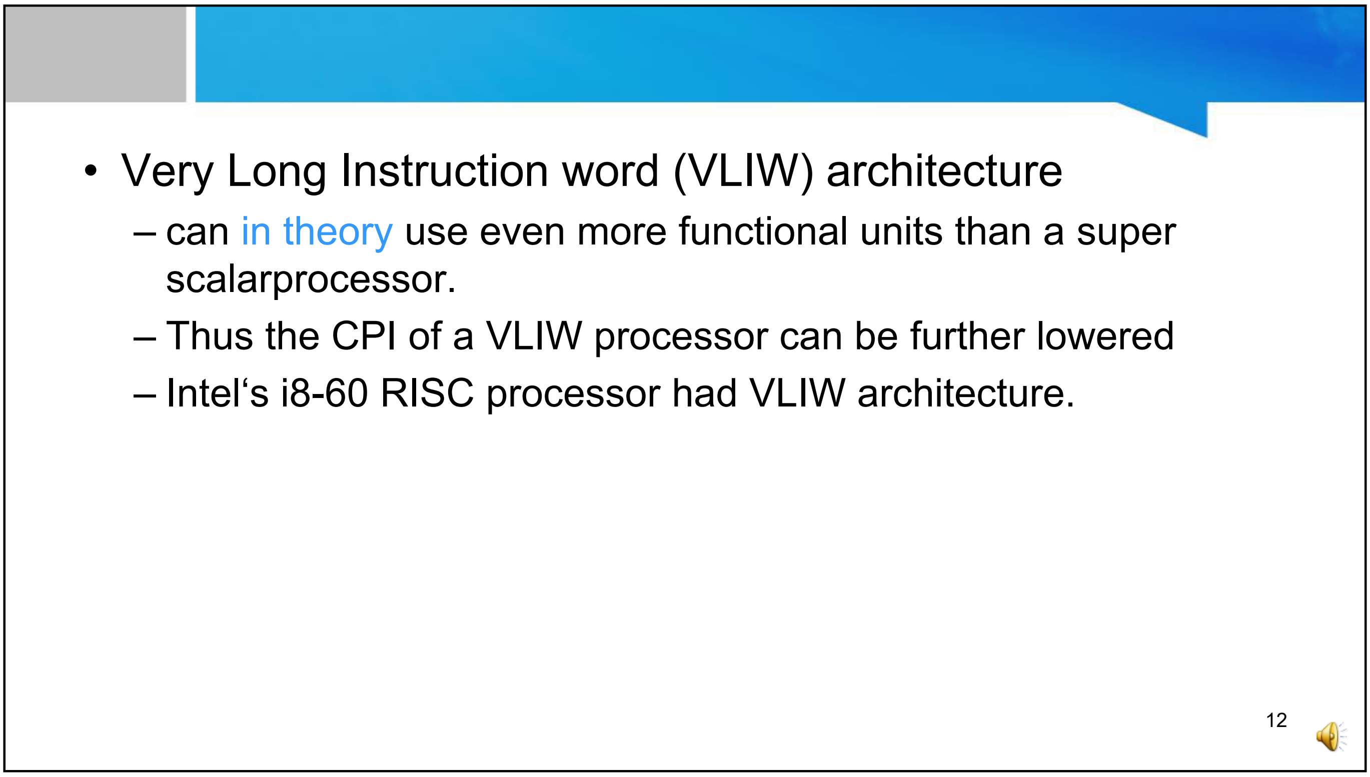

- Very Long Instruction word (VLIW) architecture
  - can in theory use even more functional units than a super scalarprocessor.
  - Thus the CPI of a VLIW processor can be further lowered
  - Intel‘s i8-60 RISC processor had VLIW architecture.

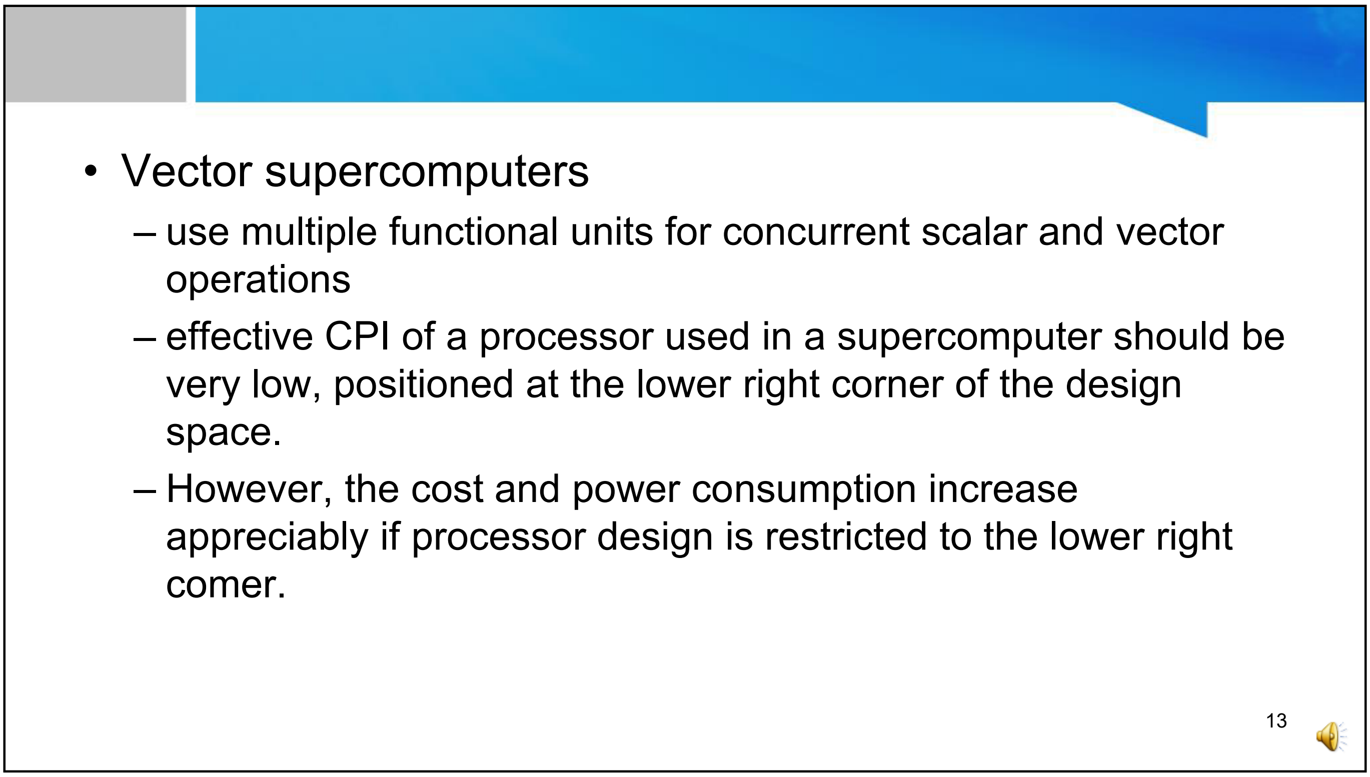

- Vector supercomputers
  - use multiple functional units for concurrent scalar and vector operations
  - effective CPI of a processor used in a supercomputer should be very low, positioned at the lower right corner of the design space.
  - However, the cost and power consumption increase appreciably if processor design is restricted to the lower right comer.

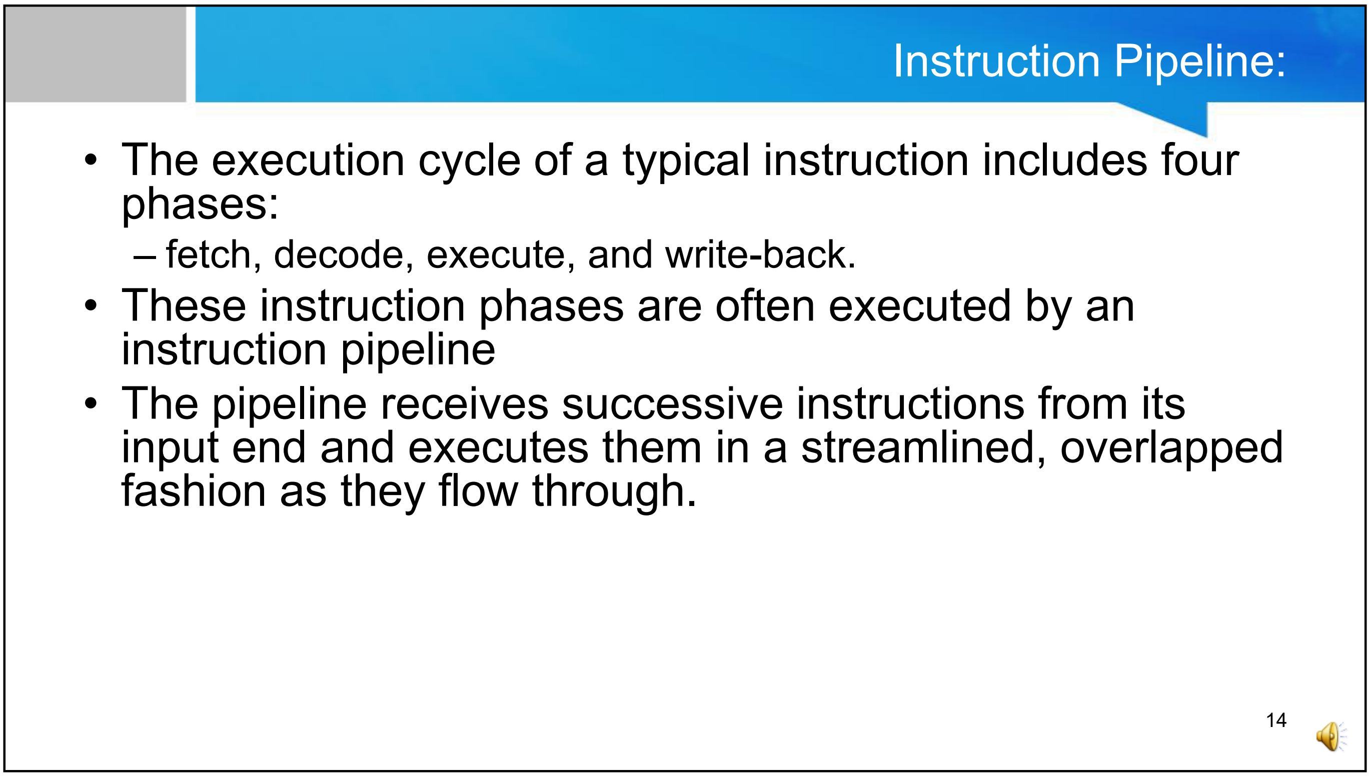

Instruction Pipeline:

- The execution cycle of a typical instruction includes four phases:
  - fetch, decode, execute, and write-back.
- These instruction phases are often executed by an instruction pipeline
- The pipeline receives successive instructions from its input end and executes them in a streamlined, overlapped fashion as they flow through.

Pipeline Definitions

- Instruction pipeline cycle
  - the time required for each phase to complete its operation (assuming equal delay in all phases)
- Instruction issue latency
  - the time (in cycles) required between the issuing of two adjacent instructions
- Instruction issue rate
  - the number of instructions issued per cycle (the *degree of a superscalar)*

- Simple operation latency
  - the delay (after the previous instruction) associated with the completion of a simple operation (e.g. integer add) as compared with that of a complex operation (e.g. divide).
- Resource conflicts
  - when two or more instructions demand use of the same functional unit(s) at the same time.

- A base scalar processor is defined as a machine with one instruction issued per cycle
- has a one-cycle latency for a simple operation
- has a one-cycle latency between instruction issues
- can be fully utilized if instructions can enter the pipeline at a rate on one per cycle
- **The effective CPI rating is 1 for the ideal pipeline**

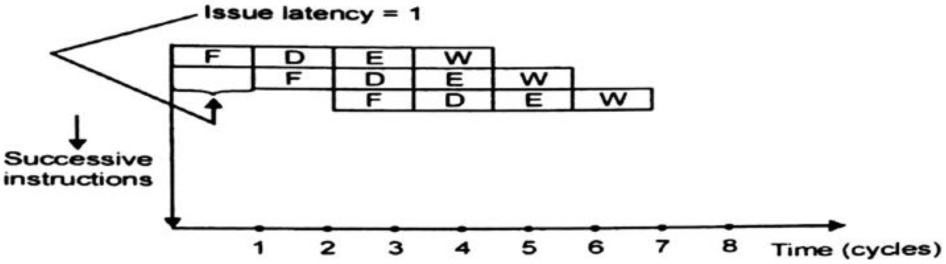

- If the instruction issue latency is two cycles per instruction the pipeline can be underutilized
- The effective CPI rating is 2

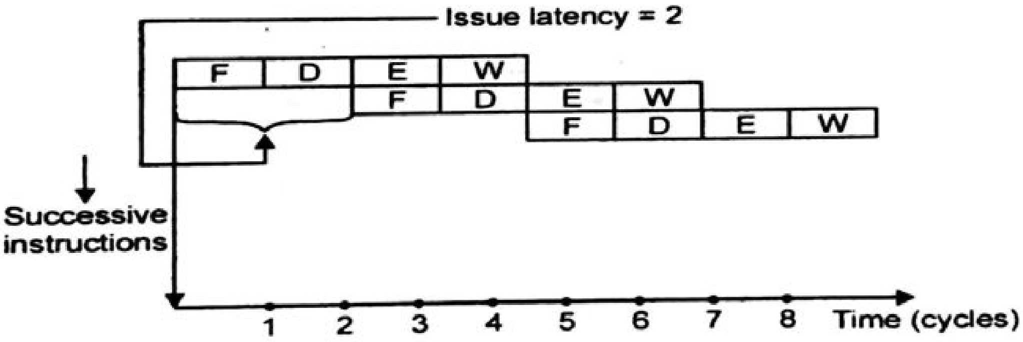

- the pipeline cycle time is doubled by combining pipeline stages.
- the fetch and decode phases are combined into one pipeline stage, and execute and write-back are combined into another stage.
- This will also result in poor pipeline utilization
- reduce the performance by one-half, compared with the ideal case

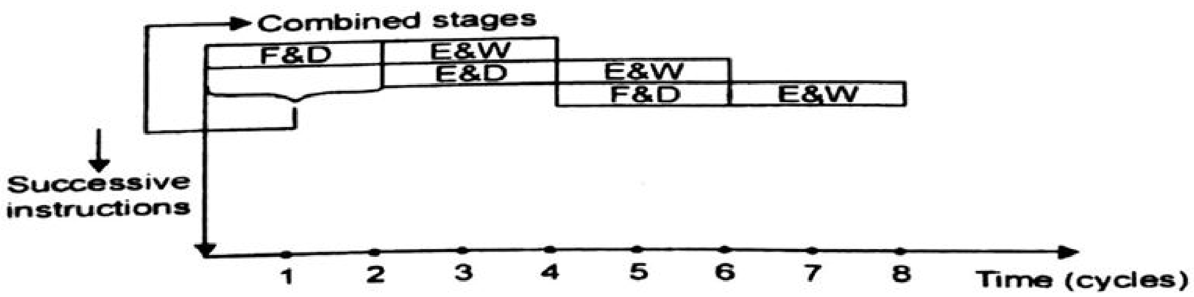

Data path architecture of a scalar processor

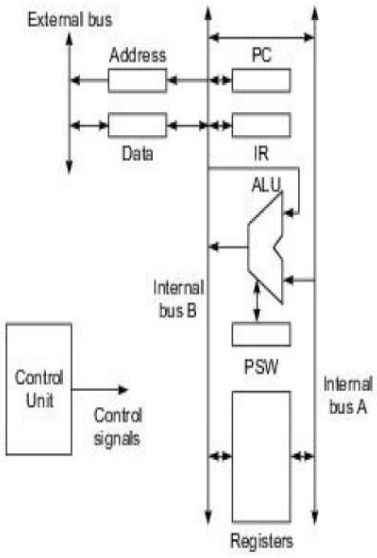

Fig. 4.3 Data path architecture and control unit of a scalar processor

- Here data path architecture and control unit of a typical, simple scalar processor without an instruction pipeline.
- Main memory, ID controllers, etc. are connected to the external bus.
- The control unit generates control signals required for the fetch, decode, ALU operation , memory access and write result phases of instruction execution.
- The control unit itself may use micro coded logic (CISC) or hardwired logic (RISC).

- Basic Scalar Processor is a machine with following features
  - 1 instruction issued per cycle
  - 1 cycle latency for simple operation
  - 1 cycle latency between instruction issues

Processors & Coprocessors

- Central processor of computer is called CPU
  - Scalar processor
  - Multiple functional units
  - Floating point accelerator
- Floating point unit can be coprocessor
  - Attached with CPU
  - Executes instructions dispatched by CPU
  - Can’t be used alone, can’t handle I/O operations

Architectural Models of a basic Scalar processor

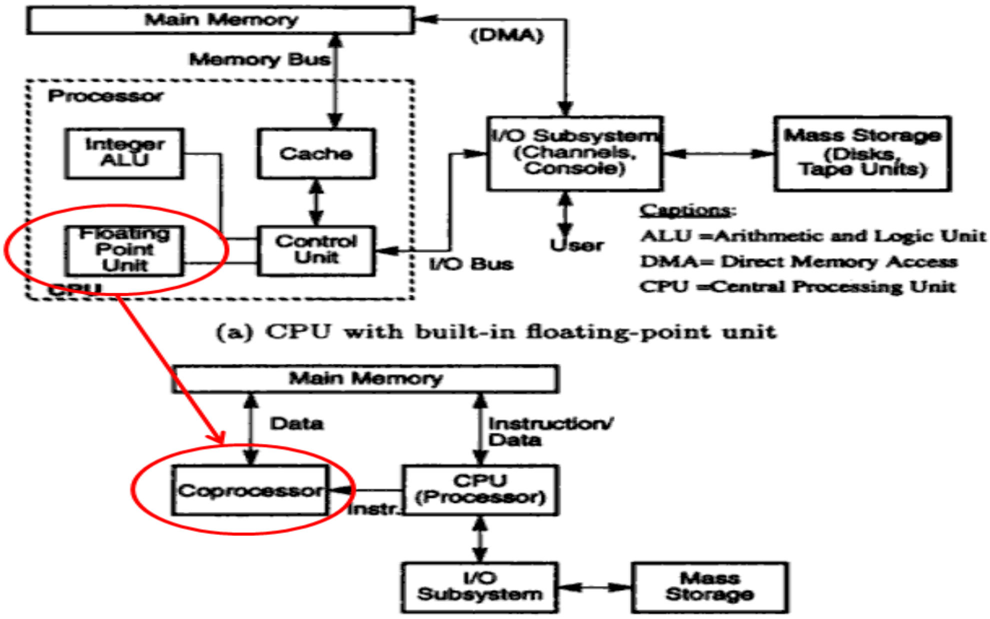

(a) CPU with built-in floating-point unit

(b) CPU with an attached coprocessor

Instruction-Set Architectures

- The instruction set of a computer specifics the primitive commands or machine instructions that a programmer can use in programming the machine.
- The complexity of an instruction set is attributed to the
  - instruction formats
  - data formats,
  - addressing modes.
  - general-purpose registers,
  - opcode specifications, and
  - flow control mechanisms used.

➢ **Examples of instruction set**

- ADD - Add two numbers together.
- COMPARE – Compare numbers.
- IN - Input information from a device, e.g.,keyboard.
- JUMP - Jump to designated RAM address.
- LOAD - Load information from RAM to the CPU.
- OUT - Output information to device, e.g.,monitor.
- STORE-Store information to RAM.

- Two types of Instruction-Set Architectures are:
  - – Complex Instruction Set Computers
  - – Reduced Instruction Set Computers

Complex Instruction Set Computing (CISC)

- HLL statements directly implemented in hardware. Add more and more functions into the hardware
- instruction set very large & complex
- Characterized by micro programmed control with Control ROM
- Typically contains 120 - 350 instructions
- Variable instruction format (16 - 64 bit)
- a few (8 - 24) general purpose registers
- Clock rate (33 - 50Mhz), CPI (2 -15)
- Lot of memory based instructions
- Unified cache design
- More than a dozen addressing modes
- Improve execution efficiency

Reduced Instruction Set Computing (RISC)

- Only 25% of large set of instructions used frequently 95% of the time. all these rare instructions to software
- reduced instruction set
- Characterized by hardwired control without Control ROM
- Typically contains less than 100 instructions
- Fixed instruction format (32 bit)
- Lot of general purpose registers (32 -192)
- Clock rate (50 - 150Mhz), CPI (< 1.5)
- Mostly register based instructions
- Split data and instruction cache design
- Only 3 – 5 addressing modes
- Memory access only by load/store instructions

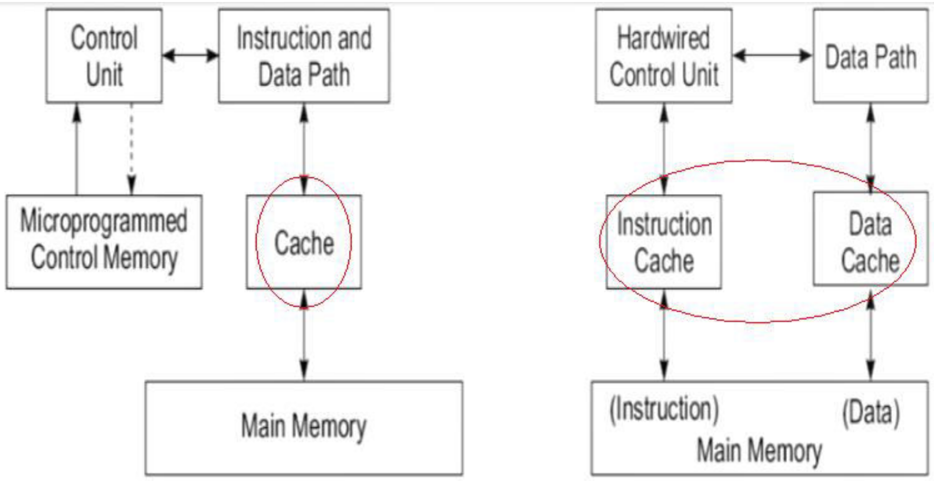

(a) The CISC architecture with microprogrammed control and unified cache

(b) The RISC architecture with hardwired control and split instruction cache and data cache

CISC Scalar Processor

❑ Scalar processor executes with scalar data

- ➢Simple models work with integer instructions using fixed point operands
- ➢Complex models work with integer and floating point operations

- Both integer unit and floating point unit may be present in same CPU
- Ideally, its performance should be that of instruction pipeline with one instruction fed per clock cycle
- Practically, it works in under pipelined situation due to data dependencies, resource conflicts, branch penalties, etc.

❑ Design Philosophy - CISC

1.Implement useful instructions in hardware, resulting in shorter program length and lower software overhead.

2.However, this is achieved at the expense of lower clock rate and higher CPI

Balance between the two required !

Representative CISC

- VAX 8600
- Motorola MC68040

VAX 8600

- Introduced by Digital Equipment Corporation
- Typical CISC architecture with Micro programmed control
- Instruction set contains 300 instructions with 20 different addressing modes
- CPU consist of two functional units
  - for execution of floating point and
  - integer instructions
- Unified cache holds both instructions and data
- 16 GPRs in instruction unit and instruction pipelining has six stages

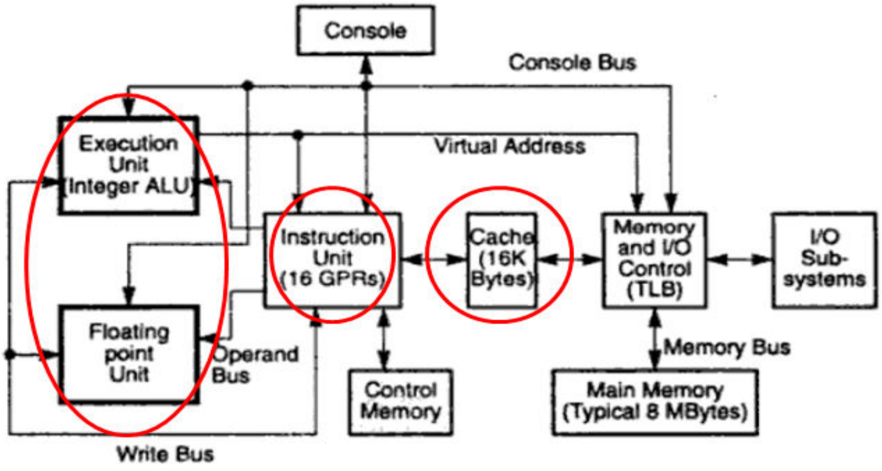

Captions:
CPU = Central Processor Unit
TLB = Translation Lookaside Buffer
GPR = General Purpose Register

The VAX 8600 CPU, a typical CISC processor architecture.

- Instruction unit
  - It prefetches and decode instructions
  - Handle branching operations
  - Supply operands to 2 functional units in a pipelined manner
- Translation look aside buffer
  - Used in memory control unit for fast generation of physical address from virtual address

Motorola MC68040

- Consist of 1.2 million transistors
- 100 instructions
- 16 GPRs
- 18 addressing modes
- Separate data and instruction cache
  - each of 4KB
- Separate MMUs supported by Address Translation Cache
  - This is equivalent to TLB
- Data format range from8 to 80 bits present in separate memory units

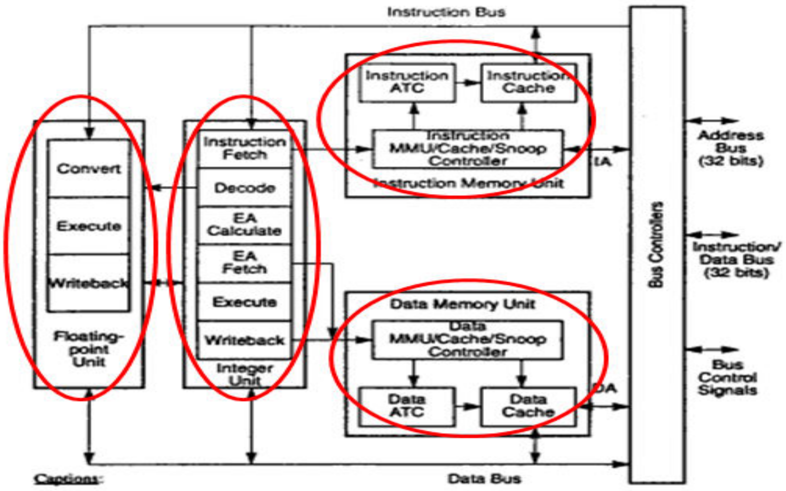

Captions:
IA = Instruction Address
DA = Data Address
EA = Effective Address
ATC = Address Translation Cache
MMU = Memory Management Unit

Architecture of the MC68040 processor

- CPU has 2 units
  - Integer unit
  - Floating point Unit
- Integer unit
  - Has 6 stage instruction pipeline
  - All instructions are decoded in this unit
  - Floating point instructions are forwarded to floating point unit
- Floating point Unit
  - Has 3 pipeline stages

- Memory units
    - Data memory units
    - Instruction memory units
  - Separate instruction and data bus used for instruction and data memory units
  - Both bus are 32 bit wide
- Each of two ATC have 64 entries
  - Provide faster translation from virtual to physical addresses

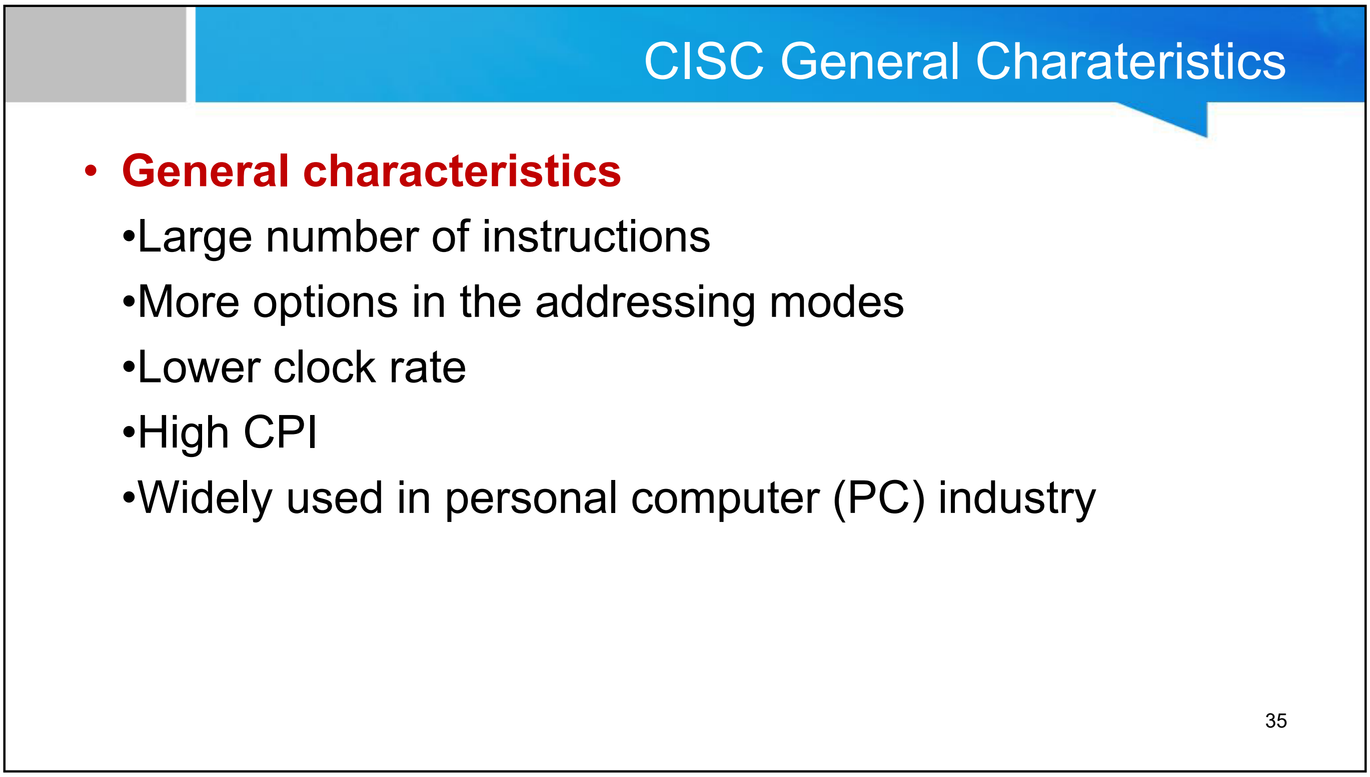

CISC General Charateristics

- **General characteristics**
  - Large number of instructions
  - More options in the addressing modes
  - Lower clock rate
  - High CPI
  - Widely used in personal computer (PC) industry

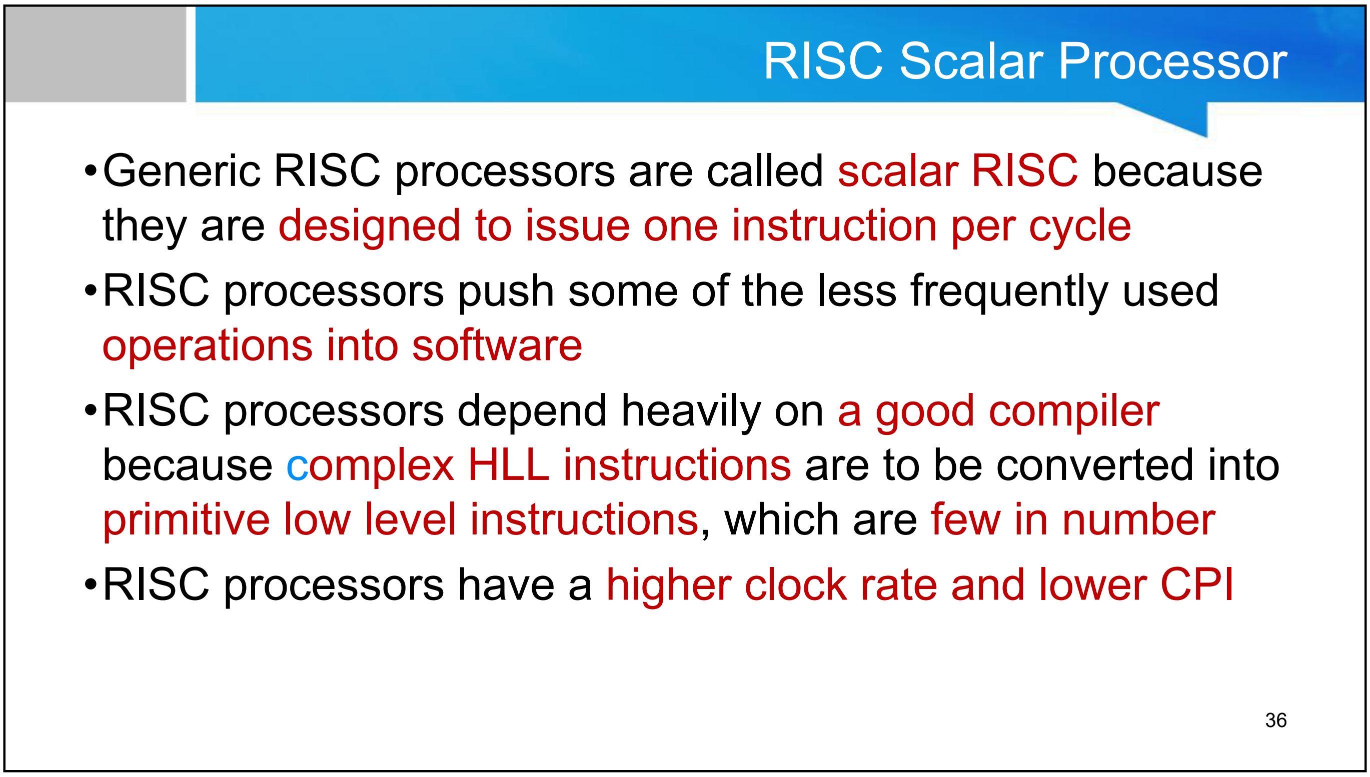

RISC Scalar Processor

- Generic RISC processors are called scalar RISC because they are designed to issue one instruction per cycle
- RISC processors push some of the less frequently used operations into software
- RISC processors depend heavily on a good compiler because complex HLL instructions are to be converted into primitive low level instructions, which are few in number
- RISC processors have a higher clock rate and lower CPI

Advantages of RISC

- **Speed:** Since a simplified instruction set allows for a pipelined, super scalar design, RISC processors often achieve 2 to 4 times the performance of CISC processor using comparable semiconductor technology and the same clock rates.
- **Simpler Hardware:** Because the instruction set of a RISC processor is so simple, it uses up much less chip space; extra functions such as memory management units or floating point arithmetic units, can also be placed on the same chip. Smaller chips allow a semiconductor manufacturer to place more parts on a single silicon wafer, which can lower the per-chip cost dramatically.
- **Shorter Design Cycle:** Since RISC processors are simpler than corresponding CISC processors, they can be designed more quickly, and can take advantage of other technological developments sooner than corresponding CISC designs, leading to greater leaps in performance between generations.

- **General characteristics**
  - All use 32-bit instructions
  - Instruction set consist of less than 100 instructions\High clock rate
  - Low CPI

SPARC

- SPARC stands for scalable processor architecture
- SPARC specification allows implementations to scale from processors required in embedded systems to processors used for servers.
- exceptionally high execution rates(MIPS) and short time-to-market development schedules.
- Scalability is due to use of number of *register windows*
- Floating point unit (FPU) is implemented on a separate chip

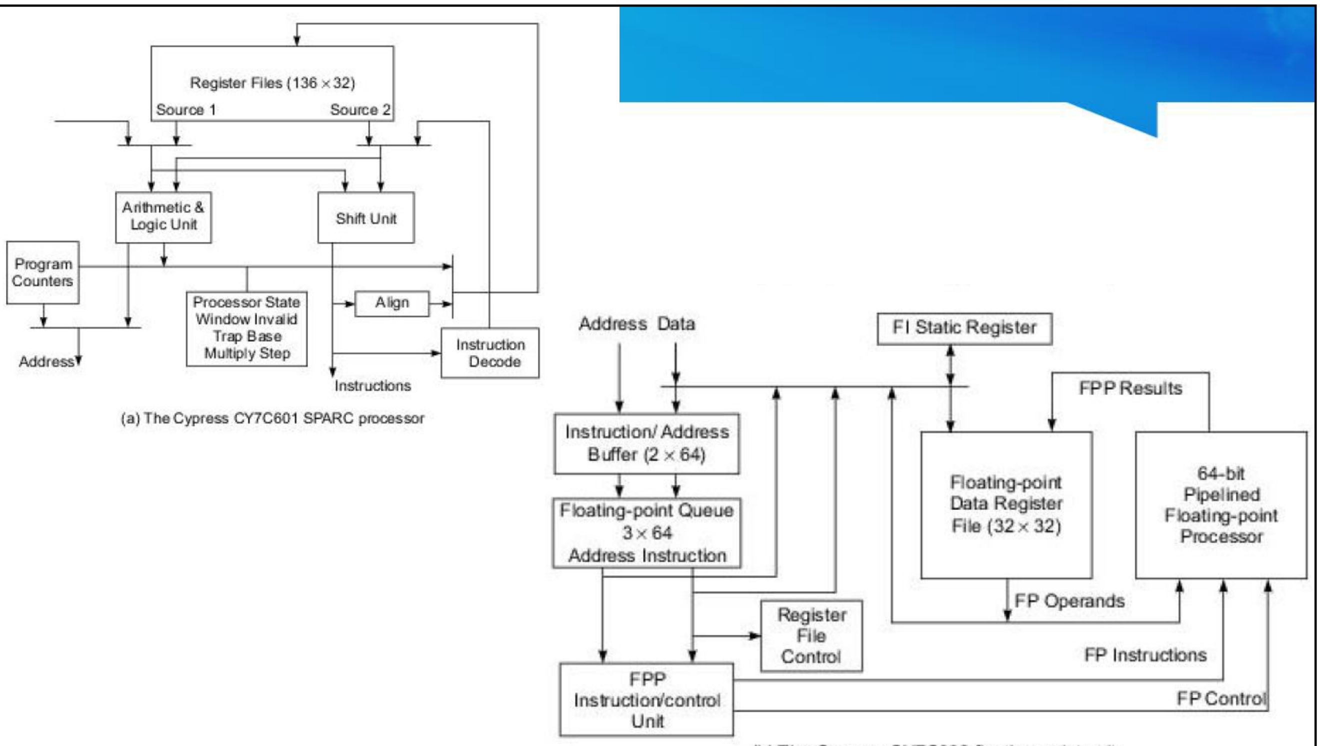
Register Files (136 × 32)
Source 1
Source 2
Arithmetic & Logic Unit
Shift Unit
Program Counters
Processor State Window Invalid Trap Base Multiply Step
Align
Instruction Decode
Address
Instructions
Address Data
FI Static Register
FPP Results
Instruction/ Address Buffer (2 × 64)
Floating-point Queue 3 × 64 Address Instruction
Floating-point Data Register File (32 × 32)
64-bit Pipelined Floating-point Processor
FP Operands
Register File Control
FP Instructions
FPP Instruction/control Unit
FP Control

(a) The Cypress CY7C601 SPARC processor

(b) The Cypress CY7C602 floating-point unit

RISC - Example 1 (Window Registers)

- ❖ SPARC runs each procedure with a set of thirty two 32-bit registers
- Eight of these registers are global registers shared by all procedures
- Remaining twenty four registers are window registers associated with only one procedure
- Concept of using overlapped registers is the most important feature introduced
- Each register window is divided into three 8 register sections – Ins, Locals and Outs
- Locals are only locally addressable by each procedure and Ins & Outs are shared among procedures

•**Input registers** : arguments are passed to a function

•**Local registers** : to store any local data.

•**Output registers** : When calling a function, the programmer puts his argument in these registers.

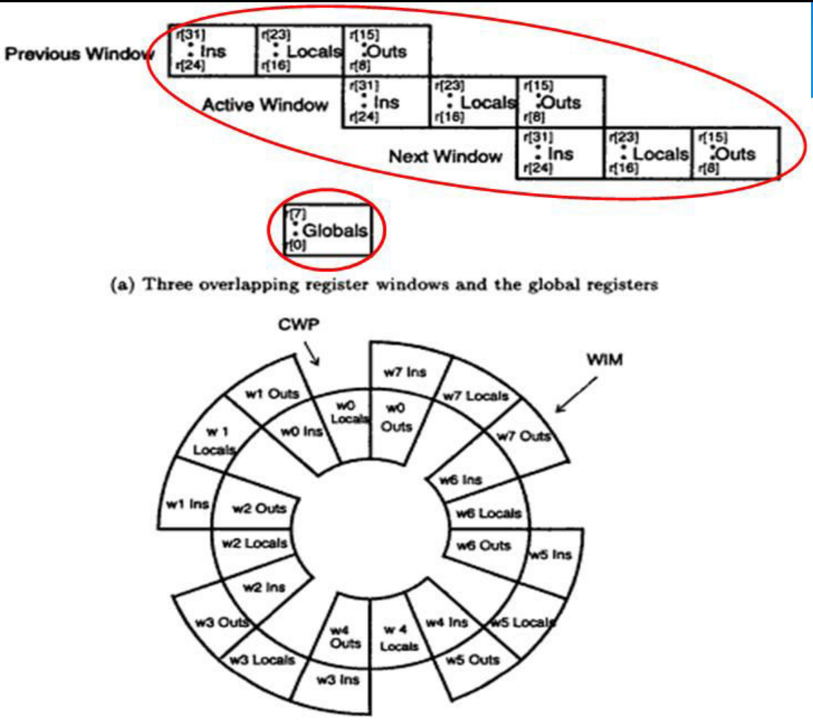

(a) Three overlapping register windows and the global registers

(b) Eight register windows forming a circular stack

RISC- Register Window

- At any time, an instruction can access the following

  8 global registers and a 24 bit register - window.

- A register window comprises a 16-register set- divided into
  - 8 in and
  - 8 local registers-
  - together with the 8 in registers of an adjacent register set, addressable from the current window as its out registers.
- When a procedure is called, the register window shifts by sixteen registers, hiding the old input registers and old local registers and making the old output registers the new input registers.
- The current (active) window into the r registers is given by the **current window pointer (CWP) register.** It is always decremented.

- **Window Invalid Mask**
  - set as 1 for oldest window…if accessed, then trap occurs, its contents saved on to stack, WIM rotated 1 bit and next lowest window set as oldest .
- **Trap base register –** pointer to trap handler
- **Special Register** to create a 64-bit product in multiple step instructions
- Overlapping windows save time in inter procedure communication, faster context switching.

RISC – Intel i860

- 64 bit RISC processor on a single chip
- It executes 82 instructions, all of them in single clock cycle
- There are nine functional units connected by multiple data
- paths
- There are two floating point units namely *multiplier unit* and *adder unit,* both of which can execute concurrently
- Dual operation add-and-multiply and subtract-and-multiply

  eg: $C = A * S2 + S1$
- Merge register used by vector integer instructions
- Graphics unit –integer operations(8, 16, 32 bit pixel data types)- 3DDrawing

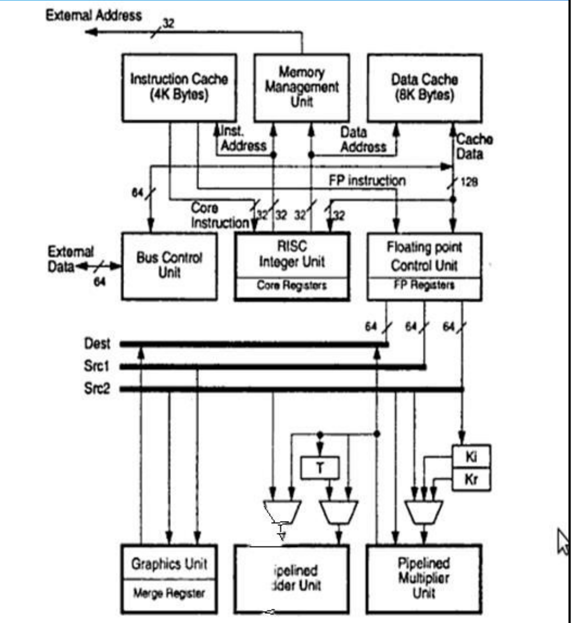

Functional units and data paths of the Intel i860 RISC microprocessor

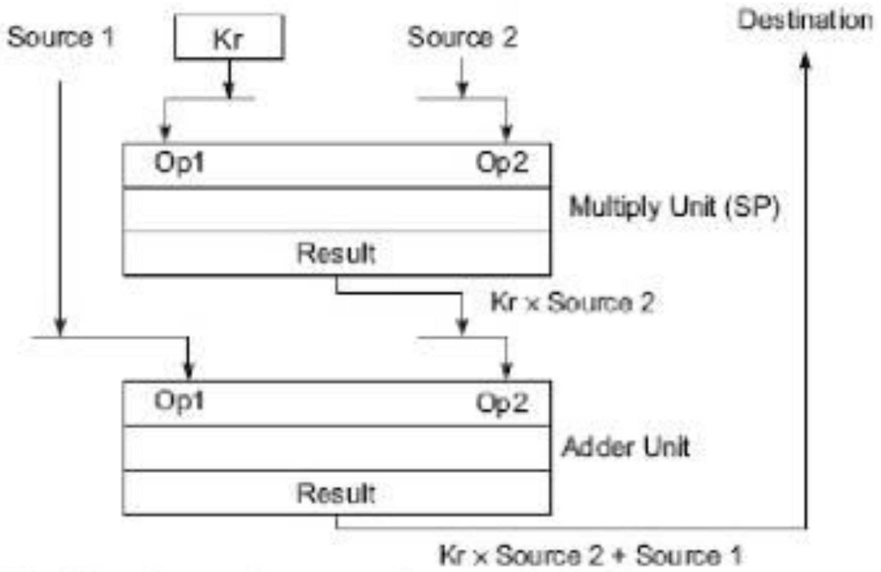

Dual floating-point operations in the i860 processor

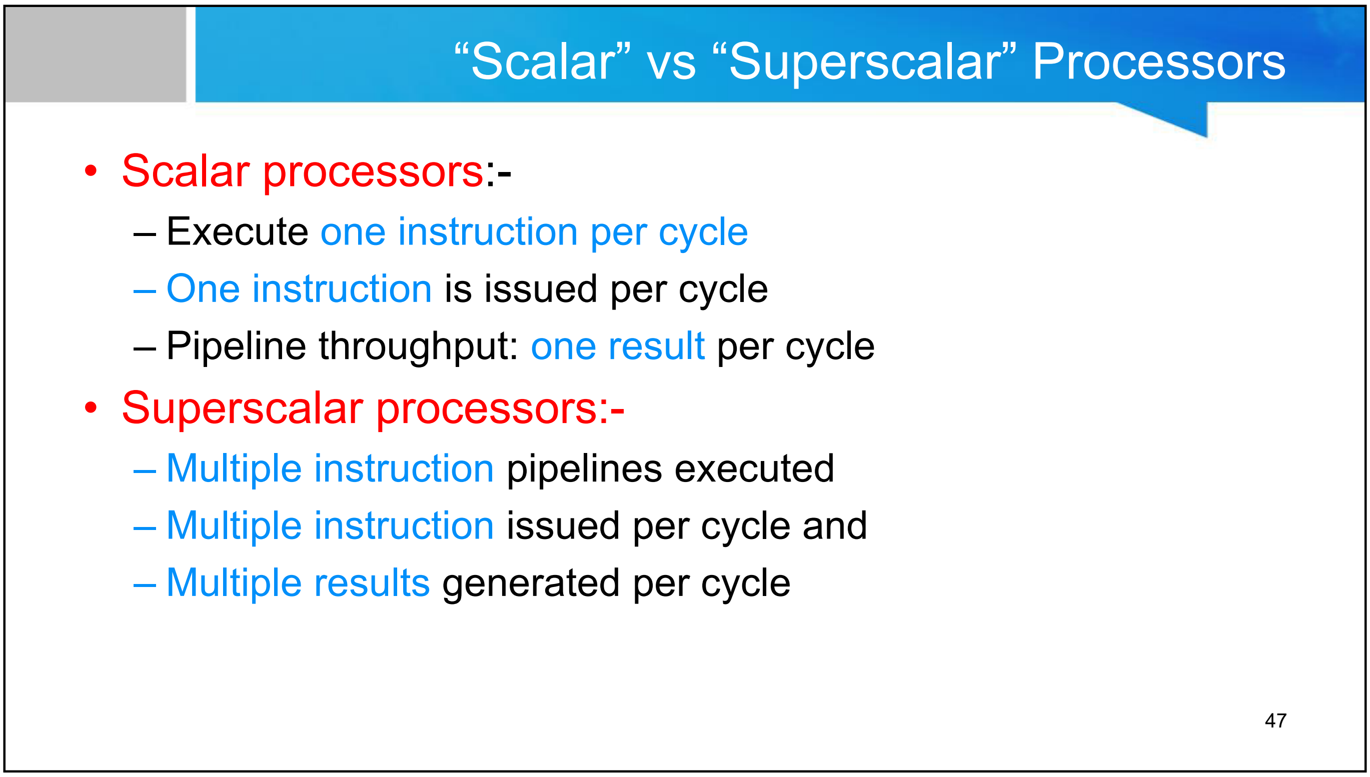

"Scalar" vs "Superscalar" Processors

- Scalar processors:-
  - Execute one instruction per cycle
  - One instruction is issued per cycle
  - Pipeline throughput: one result per cycle
- Superscalar processors:-
  - Multiple instruction pipelines executed
  - Multiple instruction issued per cycle and
  - Multiple results generated per cycle

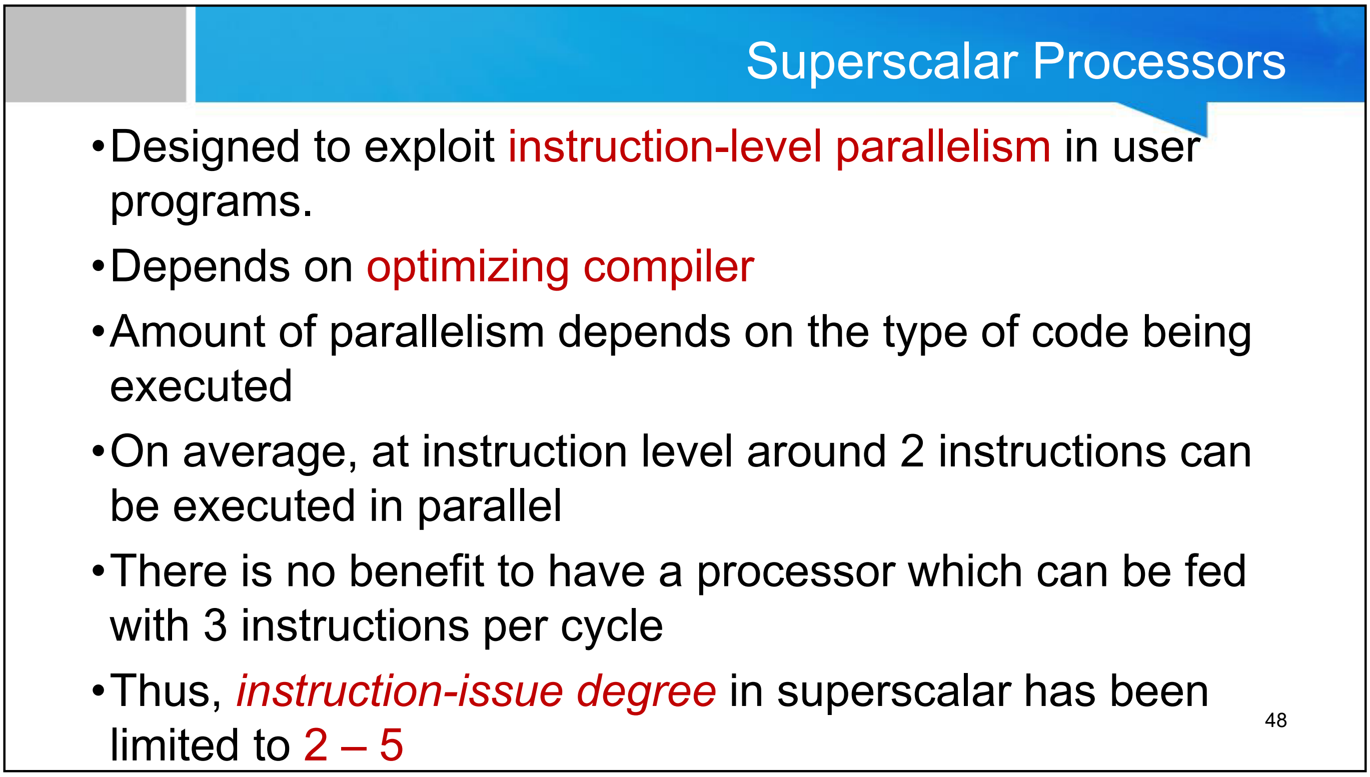

Superscalar Processors

- Designed to exploit instruction-level parallelism in user programs.
- Depends on optimizing compiler
- Amount of parallelism depends on the type of code being executed
- On average, at instruction level around 2 instructions can be executed in parallel
- There is no benefit to have a processor which can be fed with 3 instructions per cycle
- Thus, *instruction-issue degree* in superscalar has been limited to 2 – 5