

How are high speeds being realized?

- **Implicit Parallelism** Faster and faster processors combined with multiple core shared memory systems
- **Explicit Parallelism** More and more medium/coarse grain parallelism, utilizing distributed memory systems with explicit message passing
- **Implicit/Explicit Parallelism** Medium/Coarse grain parallelism facilitated through physically Shared Memory Systems

Single Core Implicit Parallelism

- Serial Parallelism, Peephole Optimizations, Pipelining
- Mostly in the **order of 2-6**
- Inherently part of processor/cache/memory design
- Requires no active involvement of the programmer (it's for free)
- Enabled through the explosion of transistor on chip (billions on a processor IC, tens of billions on a memory IC)

(Past) Trends in Single Core Architectures:

➔ Substantial Increase in clock speeds and transistor counts

- How to utilize these resources in an efficient manner.
- Current processors use these resources in multiple functional units and execute multiple instructions in the same cycle.
- The precise manner in which these instructions are selected and executed provided impressive diversity in architectures.

Pipelining and Superscalar Execution

- The speed of a pipeline is eventually limited by the slowest stage.
- For this reason, conventional processors rely on very deep pipelines (up to 20 stage pipelines in state-of-the-art Intel Core processors).
- However, in typical program traces, every 5-6th instruction is a conditional jump! This requires very accurate branch prediction.
- The penalty of a miss-prediction grows with the depth of the pipeline, since a larger number of instructions will have to be flushed.

→→→ Multiple Pipelines (Superscalar)

Superscalar Execution

Scheduling of instructions is determined by a number of factors:

- True Data Dependency: The result of one operation is an input to the next.
- Resource Dependency: Two operations require the same resource.
- Branch Dependency: Scheduling instructions across conditional branch statements cannot be done deterministically a-priori.
- Instruction Issuing Mechanisms: the scheduler looks at a large number of instructions in an instruction queue and selects appropriate number of instructions to execute concurrently based on these factors.
- In-order vs out-of-order instruction scheduling
- The complexity of this scheduler is an important constraint on superscalar processors.

Superscalar Execution: Efficiency Considerations

- Not all functional units can be kept busy at all times.
- If during a the execution of a pipeline, the pipeline is flushed, this is referred to as vertical waste.
- If during a cycle, only some of the functional units are utilized, this is referred to as horizontal waste.
- Due to limited parallelism in typical instruction traces, dependencies, or the inability of the scheduler to extract parallelism, the performance of superscalar processors is eventually limited.
- Conventional microprocessors typically support four-way superscalar execution.

Alternative: Very Long Instruction Word (VLIW) Processors

- The hardware cost and complexity of the superscalar scheduler is a major consideration in processor design.
- To address this issues, VLIW processors rely on compile time analysis to identify and bundle together instructions that can be executed concurrently.
- These instructions are packed and dispatched together, and thus the name very long instruction word.
- This concept was used with some commercial success in the Multiflow Trace machine (circa 1984).
- Variants of this concept are employed in the Intel IA64 processors.

Very Long Instruction Word (VLIW) Processors: Considerations

- Issue hardware is simpler.
- Compiler has a bigger context from which to select co-scheduled instructions.
- Compilers, however, do not have runtime information such as cache misses. Scheduling is, therefore, inherently conservative.
- Branch and memory prediction is more difficult.
- VLIW performance is highly dependent on the compiler. A number of techniques such as loop unrolling, speculative execution, branch prediction are critical.
- Typical VLIW processors are limited to 4-way to 8-way parallelism.

Limitations of Memory System Performance

- Memory system, and not processor speed, is often the bottleneck for many applications.
- Memory system performance is largely captured by two parameters: **latency** and **bandwidth**.
- Latency can be improved by providing caches between processor and memory
- Bandwidth can be improved by increasing the amount of memory interleaving (banks) and thereby increasing memory block size.

Impact of Memory Bandwidth: an Example

Consider the following code fragment, which sums columns of the matrix `b` into a vector `column_sum`:

```
for (i = 0; i < 1000; i++)
    column_sum[i] = 0.0;
    for (j = 0; j < 1000; j++)
        column_sum[i] += b[j][i];
```

- ➔ Normally the vector `column_sum` is small and easily fits into the cache.
- ➔ The matrix `b` is accessed in a column order, resulting in very bad striding behavior, reducing memory bandwidth significantly

Impact of Memory Bandwidth: an Example

We can fix the code as follows:

```
for (i = 0; i < 1000; i++)
    column_sum[i] = 0.0;
for (j = 0; j < 1000; j++)
    for (i = 0; i < 1000; i++)
        column_sum[i] += b[j][i];
```

In this case, the matrix is traversed in a row-order and performance can be expected to be significantly better.

Other ways of reducing (memory) latencies

- Multithreading allows delays to be hidden by delaying execution of one thread in favor of a thread which is not delayed.
- Prefetching allows data to be put in cache before the processor actually needs the data

Multiple Core Implicit Parallelism

- Basis was laid down in the 60's,70's and 80's
- Based on Shared Memory Architectures
- Enriched with Multiple Layers of Shared and Private/Local Cache
- Main Issue: **Fast/Parallel Shared Memory Access & Data Coherence**
- Therefore: Scalability Issues

Parallel Shared Cache/Memory Access

- For Multi Core (8/16 cores) solved by providing Memory Bus bandwidth sufficiently large, so that relatively large cache lines can be simultaneously accessed.
- For larger scale parallel platforms interconnection networks are needed, see the slides on explicit parallelism

