Performance evaluation of computer systems

J. Daniel García Sánchez (coordinator)
David Expósito Singh
Javier García Blas

Computer Architecture
ARCOS Group
Computer Science and Engineering Department
University Carlos III of Madrid

1. Module structure

This module is structured in two lessons:

1. **Trends and performance evaluation**. It introduces trends in technology, power, and energy, as well as trends regarding cost evolution. It aslo introduces performance evaluation basic concepts.

2. **Storage and reliability**. It introduces a historic perspective of storange and the concepts of reliability and availability. Besides, it combines both aspects by providing the basics of RAID storage.

2. Trends and evaluation

This lesson has the following general structure:

1. Technology trends.
2. Power and energy trends.
3. Cost trends.
4. Performance evaluation.

2.1. Technology trends

It is important to acknowledge that diverse technology changes have an impact on the processor architecture design. In particular, there are two highly relevant observations to be made.

On one hand, processor performance has improved in the past at a rate much higher than memory and storage systems. This puts much pressure in data access and justifies the addition of hierarchies to the memory systems to mitigate the problem.

On the other hand, both in processors and memory and storage systems, bandwidth and throughput has improved at a rate much higher than latency. To hide the latency problem, processor design makes extensive use of instruction level parallelism.

2.2. Power and energy trends

In general, energy can be seen as power consumed along a time interval. Consequently, the most adequate metric when comparing improvements achieved in a system is usually energy, as that is the metric considering eventual improvements in applications processing time.

In current CMOS technologies energy consumption is derived from transistor switching. Due to this, special attention is paid to the concepts of dynamic energy and dynamic power. Both depend on capacitive load and voltage, but power is the only one that also depends on switching frequency.

To reduce power and energy those observations become relevant. On one hand, a way to reduce power and energy is by decreasing operational voltage. It is for that reason that processors have moved from operation at 5 V to operate at values very near to 1 V. Another mechanism used in processor design is reducing transistors fan-out as this decreases the capacitive load.

Even though, a characteristic in processor evolution has been the increase in number of transistors as well as in frequency. this has led to moving from 2 W in Intel 80386 to 130 W in Intel Core i7. This is also one of the reasons why since around 2005 clock frequency is stagnated between 3 and 4 GHz.

2.3. Cost trends

There are multiple reasons why computer manufacturing cost, and particularly processor manufacturing cost, decreases over time.

A very important reason is derived from the learning curve principle. According to this principle performance of a manufacturing process is increased over time. As performance is the percentage of devices surviving the manufacturing process, an improvement in performance always translates into a reduction in manufacturing cost.

Another reason for cost reduction is manufacturing volume. As a general rule, the greater the manufactured volume is, the more unit cost is reduced.

Finally, when several manufacturers compete in the same product segment competence also acts as a force to decrease cost due to higher competing efforts.

2.4. Performance evaluation

The definition for performance is not unique and is context dependent. For example, for certain applications a higher performance means achieving a given task in less time (latency reduction). However, for other applications a better performance means to be able to process a larger number of transactions per unit of time (increasing throughput).

It is highly important to remark that the only metric which is really reliable to compare performance of two computer systems is based on executing real programs in both systems.

An highly used approach for computing systems evaluation is running *benchmarks*. Those are a set of programs used to evaluate performance of a system. Some approaches used to build benchmarks are the use of *kernels*, toy programs or the use of *synthetic* benchmarks. Other *benchmarks* use more realistic applications. In general, many of those *benchmarks* are specific purpose and oriented to concrete segments.

When trying to estimate speedup obtained in a system when one of its elements is improved, Amdahl's Law helps to perform such estimations if the speedup of the improved element and its percentage of participation in the system can be both estimated. this law has multiple applications, from improvements in processor designs and their impact in instructions sets to the impact of parallelization of a program segment.

A way to express processing speed is measuring cycles per instruction (CPI). The lower the average value for this metric the higher the performance exhibited by processor is when running a given application. It is very important to keep in mind that *CPI* value is usually different when running different applications. Usually, the needed time to run an application depends on its *CPI*, the number of executed instructions, and the clock cycle. On the other hand, not all the instructions require the same number of cycles to be executed. Consequently, instructions are usually aggregated by categories and a weighted average is computed to obtain the global *CPI*.

3. Storage and reliability

This lesson has the following general structure:

1. Storage.
2. Reliability and availability.
3. RAID Systems.

3.1. Storage

Magnetic disks remain being the preferred systems for secondary storage. Although in some environments they start to be partially replaced by other mass storage technologies (as *flash* memories) they are still the preferred option for corporate storage systems. In general, A magnetic disk offers a large storage capacity, and an access time depending on the sequence of storage locations from the accessed data.

Disk characteristics have evolved over time. At a first stage (1956-1970) they were devices with proprietary interfaces designed for central computers. In the 70's 5.25 inches emerged with standardized interfaces. Their mass commercialization happens in the 80's at the very same time than personal computers also emerge. Finally, the popularization of laptops and other personal devices in the 90's contributed to improvements, including successive size reductions.

3.2. Reliability and availability

Reliability is defined as the probability that a system remains correctly working without failures at a given moment in time. In general, if failure rate is constant, the system under study has a reliability with an exponential distribution function. In the materials we explain how you can determine serial and parallel systems reliability starting from individual reliabilities. Please, note that serial systems tend to decrease global reliability, as they introduce single points of failure. However, parallel systems tend to improve global reliability, as they introduce redundancy.

Availability, on the other hand, is defined as the probability that the system is working at a given instant in time, although it may have suffered maintenance operations in the past. Availability can be easily determined from the average time to failure and average time to repair.

Notice that, in general, reliability considers a time interval $[0, t]$, while availability exclusively considers time instant t.

3.3. RAID Systems

A RAID system (*Redundant Array of Inexpensive Disks*) is a system using several disks to offer a single system image for a storage system with advantages in terms of reliability. As a side effect, some configurations also offer improvements in terms of storage capacity and time performance.

Some of the most common configurations are:

- **RAID 0**: Disks with block striping. Although, they do not offer advantages in terms of fault tolerance, they improve bandwidth.
- **RAID 1**: Mirror disks. They can tolerate one disk failure as the the second disk is a copy from the first one. Besides they offer improvements in read bandwidth.
- **RAID 4**: Disks with block striping and an additional disk for information redundancy. Introduces a bottleneck for writing.
- **RAID 5**: Disks with block striping and redundancy information cyclically striped. It mitigates bottlenecks from RAID 4.