

Why modern CPUs are starving (and coping with that in PyTables)

Francesc Altet

Nowadays, every computer scientist knows that, in many occasions, their computations are going to be limited by the memory subsystem throughput and latency. This is due to the last 30-year trend that drives CPU performance to increase by a faster rate than memory buses. And this trend is going to continue in the future, forcing the scientists to design their programs in a way that minimizes memory bottlenecks so as to get the maximum performance of their processors.

In the first part of my seminar, I will talk about how system builders are designing architectures that can help to cope with this in some circumstances (basically when there is temporal or spacial data locality). Then, I will suggest using advanced compression techniques in order to allow for better performance in scenarios where current architectures are not well suited (like for instance, operating with large streaming datasets that show no temporal locality at all).

In the second part, I will introduce PyTables, a Python package for dealing with extremely large amounts of data efficiently. I'll explain how PyTables applies the techniques explained above for achieving very high performance when performing intensive input/output and computations with massive amounts of data. Also, some examples of the use of PyTables in real-world applications will be shown.

Finally, a interactive demo on PyTables will be done after the seminar for anyone who is interested.