

Where is my performance?

Atul Verma, Texas Instruments



Without a doubt, multicore processing is becoming more mainstream these days. Multicore programming or parallel programming is no longer confined to esoteric applications coded by ninja programmers. Advancements in tools and multicore programming paradigms such as OpenMP have certainly made programming simpler. However, as more and more programmers get drawn to the field of multicore, some get frustrated by the results. One question I get quite often is: Why don't I get 8 times greater performance when moving from 1 to 8 cores? This is a result of wrong expectations and some misconceptions. As much as the chip architecture is important for multicore performance, there are software considerations that are equally important. Let's dive into a few of these important factors.

Squeezing maximum performance out of the multicore architecture depends a great deal on application architecture and partitioning. The more independently multiple cores can run in parallel without the need for frequent synchronization, the better off you are in terms of performance. This aspect is termed functional partitioning. The need for synchronization between cores will actually diminish overall performance as the number of cores increase. In an ideal case different cores can be assigned to completely different functions; for example, one core can handle computations, while another executes networking tasks, while a third handles security processing and so on. Innovative chip architectures like our KeyStone multicore architecture make this easier by providing elements like Multicore Navigator and Network co-processor, which can automatically micro-schedule various processing functions across on-chip accelerators, thus relieving cores even further.

The second aspect related to multicore performance is data partitioning. Functional partitioning isn't good if independent cores are operating on the same data or on data residing in a single location. Not only will the data source become a bottleneck but also the pathways leading to the data (aka system fabric) may jam up. Therefore, it is important to partition the data and if possible locate it in different places to alleviate access bottlenecks.

Deriving maximum multicore performance depends upon both hardware and

Where is my performance?

Published on Electronic Component News (<http://www.ecnmag.com>)

software. In addition to choosing the right chip architecture, it is important to partition the application appropriately and to utilize on chip acceleration wherever possible. I invite your comments and would like to know what other unexpected behavior you have experienced when working with multicore systems? What are some of the frustrations you commonly encounter?

For more posts by TI multicore experts, visit www.ti.com/multicoresmix_ecn [1]

Source URL (retrieved on 07/10/2014 - 5:10am):

<http://www.ecnmag.com/blogs/2012/04/where-my-performance>

Links:

[1] http://www.ti.com/multicoresmix_ecn