

Teraflux: Exploiting dataflow parallelism in Teradevice Computing

Description

Dataflow parallelism is key to reach power efficiency, reliability, efficient parallel programmability, scalability, data bandwidth. In this project we proposed dataflow both at task level and inside the threads, to offload and manage accelerated codes, to localize the computation, for managing the fault information with appropriate protocols, to easily migrate code to the available/working components and to respect the power / performance / temperature / reliability envelope, to efficiently handle the parallelism and have an easy and powerful execution model, to produce a more predictable behavior.

While parallel systems have been around for many years, they were usually programmed and tuned by experts. In the future large scale systems will be widely available and therefore exploiting efficiently the available parallelism will have to be easy enough to be accessible by the common user. Traditional programming models are either not very efficient for every application (message passing) or difficult to scale (shared memory).

In order to address the programmability challenge we proposed the use of a compiler directive based model to support an underlying dataflow-based thread execution that is known to exploit well the available parallelism and to efficiently move around large amounts of data. In particular we proposed to use a model that offers dataflow scheduling of parallel execution threads. Combining multithreading with dataflow allows to exploit the available parallelism without the overheads of the original dataflow techniques. The multithreading dataflow model performed well for a number of classes of applications.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 14 jul 2024 - 22:46): <https://www.bsc.es/ca/research-and-development/projects/teraflux-exploiting-dataflow-parallelism-teradevice-computing>