

Seminar

[MSCE] Comparing strategies to reduce data-traffic on distributed memory architectures

Beyond the single-core era, Moore's Law was achieved by combining several processing units to multicore systems - be it in a homogenous or heterogeneous fashion. However, with increasing number of processing units, these architectures face scalability issues. Thus, the arising trend towards distributed or tile-based architectures seems to be a beneficial solution to satisfy the growing demand of processing power.

However, the shift towards distributed architectures often implies also a distributed memory hierarchy in order to avoid access hotspots. This arises potential data-to-task dislocalities, often referred to as "locality wall". Several trends, such as data-migration, task-migration or near-memory computing promise to increase data-to-task locality and to reduce the data traffic within the system - nowadays the most prominent energy sink in computer architectures.

The task during the seminar consists of an exhaustive literature study on that topic in order to describe, analyze and compare the most important approaches.

Contact

o.lenke@tum.de

tim.twardzik@tum.de

Advisors

Tim Twardzik