Linköping University
IDA Department of Computer and Information Sciences
PELAB
Prof. Dr. Christoph Kessler

LINKÖPING UNIVERSITY

Master Thesis Project Proposal (30hp):

Dynamic task migration for stream processing pipelines in heterogeneous distributed systems

We consider distributed soft real-time applications that process continuous data streams such as sensor data or video contents. Such applications can be computationally expensive, and are often organized for parallel processing by pipelining. In heterogeneous parallel and distributed systems, such as the IoT device / edge / cloud continuum, the different pipeline tasks can be internally parallel again and be more or less suitable for running on accelerators such as GPUs. In general, we have for each task multiple equivalent implementations to choose from. These could, for example, be provided by the programmer or generated from a high-level specification, such as SkePU code. *SkePU* (`https://skepu.github.io`) is a C++ based framework for portable high-level programming for heterogeneous parallel systems, developed in our group as an open-source effort.

Task In a previous master thesis project, a prototype framework for specifying and deploying pipelines of SkePU-defined tasks has been designed and implemented. This framework is currently being extended to also allow for implementations of pipeline tasks that are not written in SkePU.

The purpose of this thesis project is to make this framework *dynamic*, i.e., to allow for changing deployment options at runtime *while the application pipeline is processing data*, without corrupting data or having to shut down and restart the computation from the beginning. The gained flexibility will allow to react to changes in application requirements or resource availability after the initial deployment.

The prototype design and implementation of this dynamic re-deployment framework will build upon the previous prototype framework. Increased flexibility may consider, for example,

- migrating (i.e., dynamically re-mapping) a running stream-processing task from an overloaded node to a less loaded one;
- adapting a task internally by switching the SkePU back-end or its amount of hardware resources to use;
- adapting the voltage and clock frequency of an execution resource for energy-proportional computing.

More information is available from us on request.

Prerequisites TDDD56 Multicore and GPU Programming and/or TDDE65 Programming of parallel computers, TDDD25 Distributed systems, Advanced C++ (template metaprogramming), Linux. Also useful: TDDE31 Big data analytics.

Contact Christoph Kessler, Sajad Khosravi (`first.last@liu.se`)

Open thesis projects: `https://www.ida.liu.se/~chrke55/exjobb/open-exjobb-projects.shtml`