

Master Thesis Project Proposal (30hp):

Skeleton computing as a service in the edge-cloud continuum

High-level parallel programming aims to abstract challenging aspects of parallel and heterogeneous systems for non-expert programmers. *Algorithmic skeletons* is an interface approach based on computational patterns, such as *map*, *reduce*, and stencil operations. These patterns can be instantiated by providing a custom operator ("user-function"), which is then applied to a supplied dataset in parallel according to the particular pattern semantics. Skeleton programming frameworks and libraries such as SkePU implement skeletons as C++ constructs and provide "back-ends" for parallelism in multi-core CPUs, GPU accelerators, and multi-node clusters. The skeletons are typically provided as libraries, or in the case of *SkePU*, as a framework with both library and a custom compiler toolchain. The SkePU library is implemented in modern C++ and involves template metaprogramming. SkePU (<https://skepu.github.io>) is a long-term open-source effort at PELAB, Linköping University.

Task This thesis project will develop a method for deploying SkePU skeleton computations as microservices on heterogeneous parallel computing resources in the cloud or in edge computing resources for portable remote execution. This includes the specification and implementation of efficient interfaces and efficient operand data transfer, the remote deployment of a SkePU microservice with skeleton instantiation and invocation mechanisms, and the evaluation of the implementation for performance, portability, ease of use, and for security weaknesses. The project should also elaborate on suitable, remotely verifiable restrictions on user functions to be used with such services to avoid security loopholes, and implement a simple rule-based source code checker for user functions to statically verify absence of "dangerous" constructs, or at least avoid known attack patterns with high probability.

Inspiration for the service implementation can be taken from CORBA and subsequent service-oriented architectures such as web services, from MapReduce and Spark, and from a recent master thesis project (Svensson 2024) that extended SkePU for execution of stream-parallel applications in distributed systems. An experimental testbed with a number of Raspberry Pi and Jetson Nano units and GPU-accelerated servers is available for the evaluation.

A. Svensson: "SkePU Streaming: Distributed Stream Processing Pipelines for the Heterogeneous Edge-Cloud Continuum." Master thesis, IDA, Linköping University, 2024. urn:nbn:se:liu:diva-208530

Prerequisites TDDD56 Multicore and GPU Programming (mandatory), a course on Distributed Systems such as TDDD25 (mandatory), Advanced Programming in C++ (mandatory), Linux, operating systems, network programming.

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

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