The **Erlangen National High Performance Computing Center** (NHR@FAU) is looking for a

Bachelor/Master thesis student for

Extending a Parallel Performance Simulation Framework with Network Contention Models

The thesis will be hosted and supervised by the research division at Erlangen National High Performance Computing Center (NHR@FAU), which is led by Prof. Dr. Harald Köstler (Department of Computer Science, FAU).

Tasks

DisCostiC (Distributed **Cost** in **C**lusters**):** This cross-architecture simulation framework simulates the performance of massively parallel applications. The framework is essentially an automated version of analytical performance models on the full hierarchy of supercomputers, including cores, chips, nodes, networks, clusters, their individual inherent bottlenecks, and the interactions among them. It allows to simulate the behavior of parallel programs in a well-controlled, noise-free environment without requiring massive resources for computations and data transfers.

To describe the performance-relevant features of the program for simulation in an abstract way, a blueprint of the code must be expressed as a tailored domain-specific embedded language program. Static analysis facilitates the production of the application's blueprint in a semi-automatic manner. The simulated traces can be visualized using the Chromium web browser or the VAMPIR tool.

The simulation framework currently employs standard analytic performance models (i.e., Roofline and Hockney) and their enhanced model variants (i.e., ECM and LogGP family). **Within the thesis work, more advanced communication models and network contention will be added in order to make predictions more accurate for massively parallel programs.**

Within the bachelor/master thesis, the focus will be in the following areas:

- Getting familiar with the DisCostiC simulation framework and reproducing the current straightforward test cases
- Framework extension that incorporates network performance and contention models (e.g, LoGPC, LogGOPSC) to simulate the performance characteristics of highly parallel applications
- Performance assessment of the simulator predictions and their comparisons against actual near-production application runs
- Validation of simulator predictions using different networking properties and various metrics beyond runtime – data volume, communication time, other fundamental and derived metrics

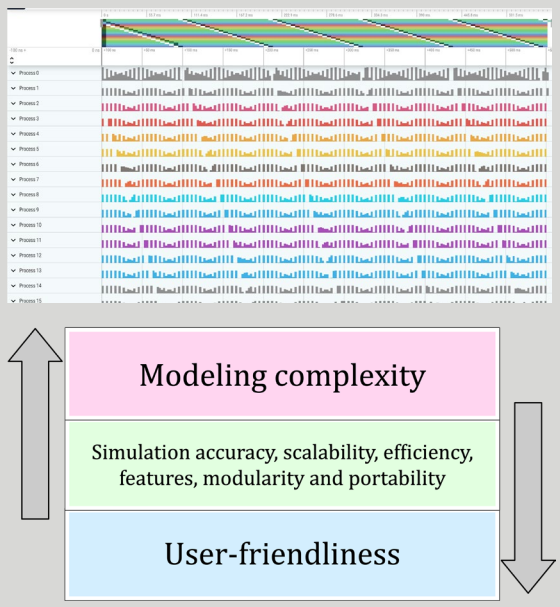

Required skills

- Student of engineering or computer science
- Profound knowledge of C/C++, Python, and the Linux OS
- Basic knowledge of code parallelization with MPI and OpenMP is preferable
- Basic knowledge of both node- and cluster-level performance engineering is preferable
- Nature of work: Theory (25%), Conception (40%), Implementation (40%)

Please direct any inquiries or applications to

Ayesha Afzal <ayesha.afzal@fau.de>
Georg Hager <georg.hager@fau.de>
Gerhard Wellein <gerhard.wellein@fau.de>