Chair for Embedded Systems

Prof. Dr. J. Henkel

Master Thesis

Multi-Agent Reinforcement Learning for Many-Core Thermal Management

The achievable performance of modern multi-core processors is mainly limited by a temperature constraint that must not be violated to avoid negative impacts on the reliability of the processor, i.e., damage due to overheating. Operating cores at lower voltage and frequency (V/f) levels may reduce the temperature, but this always comes at the cost of reduced performance. **Maximizing the performance** under a thermal constraint is only achieved by **smart decision making** to determine the optimal V/f level for each application, where high V/f levels are used only for applications that can actually benefit, and lower V/f levels are used for other applications to reduce the overall temperature.

Machine learning techniques such as **reinforcement learning** bear great potential to learn the complex application-specific dependencies of performance, power, and temperature on the V/f level.

The goal of this thesis is to develop a **multi-agent reinforcement learning** technique, where each agent is managing the cores that execute one application, with the goal of maximizing the overall performance without violating the temperature. The main challenges are how to consider the impact of agent actions on the state (temperature) of other agents, and how to achieve

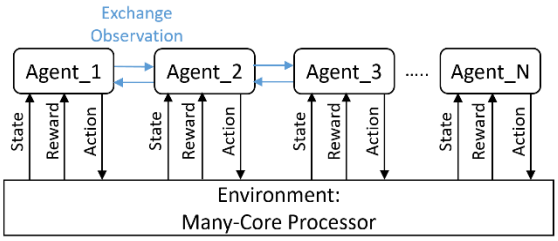

collaboration between all agents towards the shared goal of maximizing the overall performance.

Skills required for the thesis

- Programming skills (C++, Python)
- Background on machine learning
- Foreknowledge on resource management is beneficial but not required

Skills acquired within the thesis

- In depth knowledge of resource management of modern processor
- Technical writing skills
- Work in a research environment

Language

The collaboration with the colleagues would be in English.

Contact

Dr. Heba Khdr – heba.khdr@kit.edu
Dr. Martin Rapp – martin.rapp@kit.edu
Bakr Sikal – bakr.sikal@kit.edu