

Master Thesis Project Proposal (30hp):

Black-Box Models for Power Consumption Prediction of Modern Multicores

Predicting the (average) power required to perform a computation is essential for many scheduling and resource management techniques aiming at energy efficiency. However, multicore processors have become substantially more complex due to an increasing core count, heterogeneity, and chiplet-based designs with their implications on, e.g., the cache hierarchy. When it comes to predicting power consumption, decent accuracy could be obtained for earlier CPUs by exhaustive microbenchmarking feeding simple models and basic estimation techniques like least squares approximations. As an alternative to explicit power modeling, black-box approaches are conceivable where, e.g., an artificial neural network is trained using supervised learning techniques, again based on data gathered via microbenchmarking. It is of great interest whether this could be a way to accommodate growing hardware complexity and enable larger energy savings when used in scheduling.

N. Melot, C. Kessler, J. Keller, P. Eitschberger (2015). Fast Crown Scheduling Heuristics for Energy-Efficient Mapping and Scaling of Moldable Streaming Tasks on Manycore Systems. *ACM Trans. Archit. Code Optim.* 11:4. <https://doi.org/10.1145/2687653>

S. Litzinger, C. Kessler, J. Keller (2026). Implications of the amd-pstate Driver for Frequency Scaling in Energy-Efficient Task Scheduling. *PARS-Mitteilungen* 2026. <https://dl.gi.de/handle/20.500.12116/48845>

M. Colmant, R. Rouvoy, M. Kurpicz, A. Sobe, P. Felber, L. Seinturier (2018). The next 700 CPU power models. *J. Syst. Softw.* 144:382–396. <https://doi.org/10.1016/j.jss.2018.07.001>

S. Mazzola, G. Ara, T. Benz, B. Forsberg, T. Cucinotta, L. Benini (2025). Data-driven power modeling and monitoring via hardware performance counter tracking. *J. Syst. Archit.* 167:103504. <https://doi.org/10.1016/j.sysarc.2025.103504>

Task The project is intended to explore the viability of black-box power prediction approaches based on supervised learning methods for recent multicore processors that can be leveraged for scheduling techniques like *Crown Scheduling* (Melot et al. 2015). Suitable microbenchmarks will have to be curated and executed on a test system in order to collect data for training purposes. The main challenge will then lie in designing the model and experimentally identifying a concise set of relevant quantities the predictions shall be based on (e.g., operating frequency, thermals, performance counters).

Prerequisites TDDD56 Multicore and GPU Programming (mandatory), TDDE31 Big Data Analytics (recommended), TDTS08 Advanced Computer Architecture (recommended), familiarity with modeling and machine learning techniques, Linux programming skills.

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

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