Chair for Integrated Systems
TUM Department of Electrical and Computer Engineering
Technical University of Munich

Seminar

Learning Control for Predictable Latency and Low Energy

Many modern computing systems must provide reliable latency with minimal energy. Two central challenges arise when allocating system resources to meet these conflicting goals: (1) complexityÐmodern hardware exposes diverse resources with complicated interactionsÐand (2) dynamicsÐ latency must be maintained despite unpredictable changes in operating environment or input. Machine learning accurately models the latency of complex, interacting resources, but does not address system dynamics; control theory adjusts to dynamic changes, but struggles with complex resource interaction. We therefore propose CALOREE, a resource manager that learns key control parameters to meet latency requirements with minimal energy in complex, dynamic environments. CALOREE breaks resource allocation into two sub-tasks: learning how interacting resources affect speedup, and controlling speedup to meet latency requirements with minimal energy. CALOREE deines a general control systemÐ whose parameters are customized by a learning frameworkÐ while maintaining control-theoretic formal guarantees that the latency goal will be met. We test CALOREE's ability to deliver reliable latency on heterogeneous ARM big.LITTLE architectures in both single and multi-application scenarios. Compared to the best prior learning and control solutions, CALOREE reduces deadline misses by 60% and energy consumption by 13%.

Advisors

Armin Sadighi