報告番号	※甲　　　第　　　　号

主　論　文　の　要　旨

論文題目　Energy-Aware Task Allocation Exploration in Hard Real-Time Embedded Systems
(ハードリアルタイムシステムにおける低消費エネルギータスク割り当て手法)

氏　　名　　秦 洋

論 文 内 容 の 要 旨

Real-time embedded systems have become very widespread and used in industrial, medical, military and customer applications. While with the increasing demands for computing and communication, intelligent mechanisms are required to deal with the critical requirements of real-time systems, such as time constraints and power consumption. In regard to minimizing energy under time requirements, task allocation algorithms and Dynamic Voltage Frequency Scaling (DVFS) techniques have been widely explored in the industry and academia over the past decades.

Furthermore, the breakdown of Dennard Scaling causes most CPU manufacturers to focus on multicore processors as an alternative way. With the diverse needs for performance, the research on heterogeneous systems becomes more and more popular in recent decades. In terms of task allocation, the balanced workload has been proven to be the most energy-efficient for homogeneous systems, while no effective distribution method has been found to be optimal for heterogeneous systems. Moreover, the introduction of DynamIQ big.LITTLE architecture offers up configurable CPU clock domains on per-core bases, which was limited to per-cluster basis previously. It becomes increasingly difficult to propose an overall energy optimization task allocation on a multicore platform that is equipped with discrete and individual DVFS capability for each core. In this dissertation, we focus on exploring effective combinations of task allocation algorithms and non-uniform DVFS schedule policies to optimize the system energy while satisfying the time requirements on heterogeneous systems.

Our approach starts from an analytically study on a single-ISA big.LITTLE architecture which consists of two different types of cores: high-performance core and energy-efficient core. We evaluate the execution variance of different applications caused by the existence of architectural design and program characteristics, and prove that the most energy-efficient workload distribution relies on the execution variance of core type. Based on the heterogeneity of asymmetric cores, the desired workload distribution is calculated to help determine local task assignment, and an execution-variance-aware heuristic is proposed for approximating the desired distribution. We experimentally simulate randomly generated tasks and compare our method with state-of-the-art allocation solutions. Experimental results show that our approach saves more energy than the typical Little-Core-First and existing approaches with satisfying the deadline constraints of all tasks.

Next, we consider using Integer Linear Programming (ILP) formulations to solve the task allocation issue. Given the considerable time consumption, existing ILP optimizations are assumed to be effective only to the extent that the scale of the problem is small. How to use ILP to solve large allocation problems on heterogeneous multiprocessor systems to minimize energy consumption is still a challenge. In our work, we propose two ILP formulations to help find a feasible allocation. One complete ILP(1) is used to derive an intermediate solution, and the other relaxed ILP(2) is used to compute a desired minimum energy. Then the desired minimum energy can be used as a reference to evaluate the intermediate solution of ILP(1) and decide its timeout. It is verified that our proposed intermediation-based approach achieves extremely close results to the optimal solutions. Besides, to find out the best-suited platform for a given workload, a flexible design which presents flexibilities and choices in core number and core type, is considered for further energy savings. Simulation experiments of randomly generated task sets demonstrate that, compared with the fixed platform, automatically synthesizing a flexible core assignment saves more energy.

Thirdly, we shift the focus toward the intra-task DVFS as the subject of research in the task boundary of time-constrained applications for energy reduction. The problem of optimizing energy consumption with respect to intra-task DVFS scheduling can be addressed by assigning proper operational frequencies to individual basic blocks in a program while guaranteeing the deadline. Based on the profile information of a task, we first formulate the problem in terms of ILP formulations regarding different assumptions of DVFS overhead. To verify the effectiveness of ILP formulations, the most representative intra-task DVFS techniques are taken for comparisons. Experiments of randomly generated tasks and a real benchmark demonstrate that the proposed ILP method achieves greater energy savings than existing

approaches. Moreover, it determines the optimal scheduling strategy in reasonable execution time for applications with a limited number of blocks.

This dissertation presents the details of three energy optimizing scheduling approaches and evaluation experiments. Throughout the three studies, practical improvements of energy consumption over previous algorithms are achieved.