Abstract

This thesis presents fundamental work in the new area of multi-tasking on reconfigurable hardware devices (RHDs) under real-time conditions.

RHDs can now execute several hardware-tasks (computations implemented as digital circuits) in parallel due to the increasing logic capacity, as well as sequentially due to runtime reconfiguration capabilities. To use RHDs for real-time workloads of embedded system applications, scheduling techniques and execution environments that create predictable task timings are required.

Specifically, we consider the problem of scheduling periodic real-time tasks for execution on shared RHD resources, which makes the problem different from single- and multiprocessor scheduling. In our first model, tasks are modeled by their resource requirements (area), their inter-arrival period and their computation time. Tasks can be executed in parallel as long as their accumulated area does not exceed the device area and can be preempted at any time. We develop three novel *earliest deadline first (EDF)* based scheduling methods and evaluate their performance.

1. The *global EDF* scheduler assigns device resources to all active tasks globally. For this algorithm we present a *linear-time scheduling test* which proves for a given task set at design time that it will be scheduled without missing any deadlines.

2. The *partitioned EDF* scheduler divides a task set at design-time into several subsets. Each subset is scheduled independently according to the EDF rule onto separate device resources. The feasibility of the schedule is achieved by considering the single processor EDF bound during partitioning. Integer linear programming (ILP) is used to compute the optimal partitioning while a next fit heuristic computes close to optimal results for even large task sets.

3. *MSDL* is a server based scheduling method which groups tasks into periodic servers for parallel execution. At runtime, servers are loaded onto the RHD and executed according to EDF. Only one server is running at a time, making this algorithm suitable for RHDs which do not support partial reconfiguration.

In our analytic performance comparison we show that there is no dominance among the three approaches. However, our simulation studies show that, on average, global EDF outperforms partitioned EDF, which outperforms MSDL. Further we showed that the global EDF scheduling test is rather pessimistic. We analyze the reconfiguration overheads for all three approaches and include them within the schedulability tests. Simulation experiments show that *global EDF* suffers significantly more from this overhead than the other schedulers.

Two extensions of our task model are presented in order to consider the specific characteristics of reconfigurable applications: First we consider tasks for which not only one but several alternative implementations (circuits) exist. We present an ILP model to select the optimal implementations for the partitioned EDF scheduler, which considerably improves the scheduling performance. Second, we consider tasks memory usage under real-time constraints. We present a method for efficiently sharing RAM-banks without jeopardizing task deadlines, thus reducing the required memory resources.

Finally we describe our prototype FPGA based real-time kernel which, in contrast to others, is a hardware-only system with the operating system functions entirely implemented in hardware. Our prototype implements MSDL, uses full device reconfiguration, and is thus suitable for most of today's reconfigurable devices.

In summary, the thesis presents a foundation for real-time multitasking on RHDs, including models, algorithms, analysis and prototypical implementation.