Ph.D. in Information Technology: Thesis Defense

October 27th, 2021
online by Teams – at 11.00

Luca CREMONA – XXXIII Cycle

Online power modeling, monitoring and optimization for mobile computing infrastructures

Supervisor: Prof. **William Fornaciari**

Abstract:

The Internet-of-Things (IoT) revolution fueled new challenges and op- portunities to achieve computational efficiency goals. Embedded devices are required to execute multiple applications for which a suitable distribution of the computing power must be adapted at run-time. Such complex hardware platforms have to sustain the continuous acquisition and processing of data under severe energy budget constraints, since most of them are battery powered. The state-of-the-art offers several ad- hoc contributions to selectively optimize the performance considering aspects like energy, power, thermal or reliability. In this scenario, the use of hardware-level online power monitors is crucial to support the runtime power optimizations required to meet the ever increasing demand for energy efficiency. To be effective and to deal with the time-to-market pressure, the presence of such requirements must be considered even during the design of the power monitoring infrastructure. This thesis presents a power model identification and implementation strategy with two main advantages over the state-of-the-art. First, the proposed solution trades the accuracy of the power model with the amount of resources allocated to the power monitoring infrastructure. Second, the use of an automatic power model instrumentation strategy ensures a timely implementation of the power monitor regardless the complexity of the target computing platforms. To assess the effectiveness of the proposed solution the identified power monitor has been adopted to feed a power optimization scheme, based on a control theory based PID controller. Both the single-core and multi-core scenarios have been taken into consideration. The online power monitor has been validated against 8 accelerators generated through a High-Level-Synthesis flow and by considering a more complex RISC-V embedded computing platform. The all-digital power optimization scheme has been validated against the *nu+* processor, a 16-ways SIMD processor with a configurable number of cores. For the assessment of the proposed control scheme, this thesis

considers the four core configuration, running 20 applications from the WCET benchmark suite. For what concerns the power monitor, depending on the imposed user-defined constraints and with respect to the unconstrained power monitoring state-of-the-art solutions, the proposed methodology shows a resource saving between 37.3% and 81% while the maximum average accuracy loss stays within 5%, i.e., using the aggressive 20us temporal resolution. However, by varying the temporal resolution closer to the value proposed in the state of the art, i.e. in the range of hundreds of microseconds, the average accuracy loss of the power monitors is lower than 1% with almost the same overheads. In addition, the presented solution demonstrated the possibility of delivering a resource constrained power monitor employing a 20us temporal resolution, i.e., far higher the one used by current state-of-the-art solutions. The power optimization scheme, instead, shows an overhead limited to 0.86%(FFs) and 5.3%(LUTs) of the FPGA chip. The performance results are analyzed considering three quality metrics. First, the efficiency in exploiting the imposed budget (EFF_g) that is on average 98.27%. Second, the overflow of the actual average power consumption with respect to the assigned budget ($OV\ F_g$), which is limited to 1.43 mW on average. Last, the performance utility loss due to the control scheme that is limited to 1.87% on average.

PhD Committee
Prof. **Giovanni Agosta**, DEIB
Prof. **Luciano Lavagno**, Universidad Politecnica de Madrid
Prof. **Mauro Olivieri**, Universita' di Roma La Sapienza