Preface

Embedded systems are becoming more and more ubiquitous in our daily lives. Since the embedded systems interact with the physical environment, they often have to behave within specified time limits. We call such systems real-time systems. Violating the timing constraints may lead to disastrous consequences. Therefore, it must be guaranteed that the time constraints are satisfied in all situations. This is an extremely difficult task, as the number of possible system states is overwhelmingly large or even infinite.

In this book, we introduce some new methods to design and analyze real-time systems. Timing analysis of real-time systems is typically performed in a bottom-up manner, starting with the system's smallest components and then gradually up toward analyzing the system as a whole. The analysis on the program level aims to give each piece of program an upper limit of its execution time. On the component level (e.g., in a processor or a communication channel), many computation/communication tasks compete for the same platform, and the analysis should guarantee each of them obtains enough resource to process in time. On the system level, the analysis takes into account the interaction between the computation and communication activities that are distributed on various components. This book covers topics on each of these three levels.

On the program level (Part I) we study the worst-case execution time (WCET) analysis problem in the presence of caches with two commonly used replacement policies MRU and FIFO. Most of the research done with respect to the impact of caches to WCET of programs assume LRU caches. However, LRU is actually not commonly used in commercial processors because it requires a relatively complex hardware implementation. Hardware manufacturers tend to design the caches that are not using the LRU policy, particularly for embedded systems which are subject to strict restrictions in cost, power consumption, and heat. It has been found that existing methods of analysis, which is based on a qualitative classification of hits and misses in the cache, are not suitable to analyze MRU or FIFO cache. Our main contribution in this part is the development of a quantitative analysis methodology which is better suited to analyze the MRU and FIFO.

On the processor level (Part II) we address the challenges that arise due to the trend of multi-core processors. We first study some fundamental problems in multiprocessor real-time scheduling theory. Multiprocessor scheduling can be divided into two main paradigms, global and partitioned scheduling. In global scheduling, it is generally unknown what corresponds to the worst combination of contemporary software activations, which makes the analysis of such scheduling more difficult than its counterpart on single processors. Analysis of partitioned multiprocessor scheduling is simpler, but there is instead the challenge of how to allocate workload to individual processors for high resource utilization. We have studied the above theoretical problems and greatly advanced the state of the art. Moreover, we also consider practical aspects of real-time scheduling on multi-cores, to solve the interference arising between processor cores because of shared caches. We propose to use cache-aware scheduling where page-coloring are used to provide predictable cache performance of individual tasks.

On the system level, we introduce new techniques that solve the efficiency problem in the widely used real-time calculus (RTC) framework. Operations within the RTC generate curves with a frequency that is the product of the periodicity of the input curves. Therefore, RTC has an exponential complexity, and in practice the efficiency can be very low for complex systems. In this book, we present finitary real-time calculus to solve the above problems. The idea is to only maintain operations on the prefix of each trace which can affect the final result of the analysis. In that way we can show that analysis complexity is reduced to be pseudo-polynomial, and in practice the analysis efficiency is dramatically improved comparing with the ordinary RTC. An additional contribution to the RTC in this book is new analysis techniques of the earliest deadline first (EDF) scheduling algorithm in the RTC framework.

This book is the collection of some of the work I did during my Ph.D. study in Uppsala University, Sweden. I would like to thank my supervisor Wang Yi. I am truly grateful for his guidance, support, inspiration, patience, and optimism, especially when I was depressed by the negative aspects of my work as he consistently believes that there will be a way out and encourages me to carry on. I learned a lot more than computer science from him during the past 5 years. Without him, the work in this book would not have been possible to be finished (not possible to be started actually).

Kowloon, Hong Kong
September 2015

Nan Guan

http://www.springer.com/978-3-319-27196-5

Techniques for Building Timing-Predictable Embedded Systems
Guan, N.
2016, XIV, 235 p. 54 illus., 20 illus. in color., Hardcover
ISBN: 978-3-319-27196-5