

Preface

In the context of computer systems, scheduling theory is concerned with the efficient allocation of computational resources, which may be available in limited amounts, amongst competing demands in order to optimize specified objectives. Real-time scheduling theory deals with resource allocation in real-time computer systems, which are computer systems in which certain computations have a timing correctness requirement in addition to a functional one—the correct value must be computed, at the right time.

Computer systems, real-time and otherwise, are increasingly coming to be implemented upon multiprocessor and multicore platforms. The real-time systems research community has certainly recognized this reality, and has responded with enthusiasm and gusto: A large body of research has been performed addressing the various issues, challenges, and opportunities arising from this move towards multiprocessor platforms. A substantial fraction of this body of work is devoted to the scheduling-theoretic analysis of multiprocessor real-time systems.

We started out with the objective of distilling the most relevant ideas, techniques, methodologies, and results from this body of work into a foundational intellectual core for the discipline of multiprocessor real-time scheduling theory. However, the process of identifying such a core proved to be very challenging and trying. We soon discovered that we could not possibly provide comprehensive coverage of all the important ideas; indeed the sheer volume of excellent research that has been produced by the community meant that we could not even guarantee to at least mention every good idea that we came across. We have therefore instead chosen to select a self-contained collection of topics from the vast body of research literature on multiprocessor real-time scheduling theory, and to provide a cohesive, relatively deep, and complete coverage of this collection of topics. Our choice of topics was governed by (i) their relevance to the discipline; (ii) the depth of knowledge and understanding that has been acquired regarding them; and (iii) their cohesiveness—i.e., that all taken together they should comprise a complete narrative for a substantial and important subset of multiprocessor real-time scheduling theory. With these goals in mind,

we chose to focus the coverage in the book upon the real-time scheduling of sporadic¹ task systems upon preemptive identical multiprocessor platforms, generally describing positive results (such as efficient algorithms for analysis) in preference to negative ones (lower bounds, impossibility proofs, etc.)

Organization

Broadly speaking, the chapters of this book are organized into the following four categories.

1. *Background material.* Chapters 1–3 provide background and context, and define many of the terms used during the remainder of the book. They additionally serve to demarcate the scope of this book, and explain the reasoning that guided our choice of this particular scope. Chapter 4 provides a brief review of some uniprocessor scheduling theory results that will be used during our study of multiprocessor scheduling.
2. *Liu and Layland systems.* The Liu and Layland task model (described in detail in Chap. 2) is a simple formal model for representing real-time workloads that is widely used, and has been very widely studied. Chapters 5–9 provide a fairly detailed coverage of a selection of topics dealing with the scheduling and schedulability analysis of Liu and Layland task systems.
3. *Three-parameter sporadic task systems.* This generalization to the Liu and Layland task model (also described in detail in Chap. 2) has been the focus of a lot of the research in multiprocessor real-time scheduling theory. A substantial portion of this book—Chaps. 10–20—is devoted to discussing the multiprocessor scheduling of real-time systems that are represented using the three-parameter sporadic task model.
4. *Emerging topics.* Many additional topics have come to be recognized as being extremely important and relevant to the development of a comprehensive theory of multiprocessor real-time scheduling. Some preliminary research has been conducted upon these topics and a few results obtained, but the state of our knowledge is not as deep or as comprehensive for the topics discussed in Chaps. 1–20. We briefly review what is known about a few such topics towards the end of the book. Chapter 21 discusses the sporadic DAG tasks model. Chapter 22 discusses real-time scheduling upon heterogeneous multiprocessor platforms.

For further information. As stated above, we have chosen to cover only a selection of the wide range of topics that have been studied in multiprocessor real-time scheduling theory. To get a more complete sense of the scope of the discipline, we recommend that the reader browse through the proceedings from recent years of conferences such as the IEEE Real-Time Systems Symposium (RTSS), and the EuroMicro Conference on Real-Time Systems (ECRTS).

¹ We will formally define all these terms—sporadic, preemptive, identical, and efficient—in Chap. 1.

Multiprocessor Scheduling for Real-Time Systems

Baruah, S.; Bertogna, M.; Buttazzo, G.

2015, XV, 228 p. 45 illus., 4 illus. in color., Hardcover

ISBN: 978-3-319-08695-8