Preface

Preface to the First Edition

Sequencing and scheduling is a form of decision-making that plays a crucial role in manufacturing and service industries. In the current competitive environment effective sequencing and scheduling has become a necessity for survival in the market-place. Companies have to meet shipping dates that have been committed to customers, as failure to do so may result in a significant loss of goodwill. They also have to schedule activities in such a way as to use the resources available in an efficient manner.

Scheduling began to be taken seriously in manufacturing at the beginning of this century with the work of Henry Gantt and other pioneers. However, it took many years for the first scheduling publications to appear in the industrial engineering and operations research literature. Some of the first publications appeared in Naval Research Logistics Quarterly in the early fifties and contained results by W.E. Smith, S.M. Johnson and J.R. Jackson. During the sixties a significant amount of work was done on dynamic programming and integer programming formulations of scheduling problems. After Richard Karp's famous paper on complexity theory, the research in the seventies focused mainly on the complexity hierarchy of scheduling problems. In the eighties several different directions were pursued in academia and industry with an increasing amount of attention paid to stochastic scheduling problems. Also, as personal computers started to permeate manufacturing facilities, scheduling systems were being developed for the generation of usable schedules in practice. This system design and development was, and is, being done by computer scientists, operations researchers and industrial engineers.

This book is the result of the development of courses in scheduling theory and applications at Columbia University. The book deals primarily with machine scheduling models. The first part covers deterministic models and the second part stochastic models. The third and final part deals with applications. In this last part scheduling problems in practice are discussed and the relevance of the theory to the real world is examined. From this examination it becomes

clear that the advances in scheduling theory have had only a limited impact on scheduling problems in practice. Hopefully there will be in a couple of years a second edition in which the applications part will be expanded, showing a stronger connection with the more theoretical parts of the text.

This book has benefited from careful reading by numerous people. Reha Uzsoy and Alan Scheller Wolf went through the manuscript with a fine tooth comb. Len Adler, Sid Browne, Xiuli Chao, Paul Glasserman, Chung-Yee Lee, Young-Hoon Lee, Joseph Leung, Elizabeth Leventhal, Rajesh Sah, Paul Shapiro, Jim Thompson, Barry Wolf, and the hundreds of students who had to take the (required) scheduling courses at Columbia provided many helpful comments which improved the manuscript.

The author is grateful to the National Science Foundation for its continued summer support, which made it possible to complete this project.

Michael Pinedo
New York, 1994.

Preface to the Second Edition

The book has been extended in a meaningful way. Five chapters have been added. In the deterministic part it is the treatment of the single machine, the job shop and the open shop that have been expanded considerably. In the stochastic part a completely new chapter focuses on single machine scheduling with release dates. This chapter has been included because of multiple requests from instructors who wanted to see a connection between stochastic scheduling and priority queues. This chapter establishes such a link. The applications part, Part III, has been expanded the most. Instead of a single chapter on general purpose procedures, there are now two chapters. The second chapter covers various techniques that are relatively new and that have started to receive a fair amount of attention over the last couple of years. There is also an additional chapter on the design and development of scheduling systems. This chapter focuses on rescheduling, learning mechanisms, and so on. The chapter with the examples of systems implementations is completely new. All systems described are of recent vintage. The last chapter contains a discussion on research topics that could become of interest in the next couple of years.

The book has a website:

`http://www.stern.nyu.edu/~mpinedo`

The intention is to keep the site as up-to-date as possible, including links to other sites that are potentially useful to instructors as well as students.

Many instructors who have used the book over the last couple of years have sent very useful comments and suggestions. Almost all of these comments have led to improvements in the manuscript.

Reha Uzsoy, as usual, went with a fine tooth comb through the manuscript. Salah Elmaghraby, John Fowler, Celia Glass, Chung-Yee Lee, Sigrid Knust,

Joseph Leung, Chris Potts, Levent Tuncel, Amy Ward, and Guochuan Zhang all made comments that led to substantial improvements.

A number of students, including Gabriel Adei, Yo Huh, Maher Lahmar, Sonia Leach, Michele Pfund, Edgar Possani, and Aysegul Toptal, have pointed out various errors in the original manuscript.

Without the help of a number of people from industry, it would not have been possible to produce a meaningful chapter on industrial implementations. Thanks are due to Heinrich Braun and Stephan Kreipl of SAP, Rama Akkiraju of IBM, Margie Bell of i2, Emanuela Rusconi and Fabio Tiozzo of Cybertec, and Paul Bender of SynQuest.

Michael Pinedo
New York, 2001.

Preface to the Third Edition

The basic structure of the book has not been changed in this new edition. The book still consists of three parts and a string of Appendixes. However, several chapters have been extended in a meaningful way, covering additional topics that have become recently of interest. Some of the new topics are more methodological, whereas others represent new classes of models.

The more methodological aspects that are receiving more attention include Polynomial Time Approximation Schemes (PTAS) and Constraint Programming. These extensions involve new material in the regular chapters as well as in the Appendixes. Since the field of online scheduling has received an enormous amount of attention in recent years, a section focusing on online scheduling has been added to the chapter on parallel machine scheduling.

Two new classes of models are introduced in the chapter on more advanced single machine scheduling, namely single machine scheduling with batch processing and single machine scheduling with job families.

Of course, as in any new edition, the chapter that describes implementations and applications had to be revamped and made up-to-date. That has happened here as well. Two new software systems have been introduced, namely a system that is currently being implemented at AMD (Advanced Micro Devices) and a generic system developed by Taylor Software.

For the first time, a CD-ROM has been included with the book. The CD-ROM contains various sets of power point slides, minicases provided by companies, the LEKIN Scheduling system, and two movies. The power point slides were developed by Julius Atlason (when he taught a scheduling course at the University of Michigan-Ann Arbor), Johann Hurink (from the University of Twente in Holland), Rakesh Nagi (from the State University of New York at Buffalo), Uwe Schwiegelshohn (from the University of Dortmund in Germany), Natalia Shakhlevich (from the University of Leeds in England).

A website will be maintained for this book at

`http://www.stern.nyu.edu/~mpinedo`

The intention is to keep this website as up-to-date as possible, including links to other sites that are potentially useful to instructors as well as to students.

A hardcopy of a solutions manual is available from the author for instructors who adopt the book. The solutions provided in this manual have been prepared by Clifford Stein (Columbia University), Julias Atlason (Michigan), Jim Geelen (Waterloo), Natalia Shakhlevich (Leeds), Levent Tuncel (Waterloo), and Martin Savelsbergh (Georgia Tech).

I am very grateful to a number of colleagues and students in academia who have gone over the new sections and have provided some very useful comments, namely Alessandro Agnetis (Siena), Ionut Aron (T.J. Watson Research Laboratories, IBM), Dirk Briskhorn (Kiel), John Fowler (Arizona), Jim Geelen (Waterloo), Johann Hurink (TU Twente, the Netherlands), Detlef Pabst (AMD), Gianluca de Pascale (Siena, Italy), Jacob Jan Paulus (TU Twente, the Netherlands), Jiri Sgall (Charles University, Prague), and Gerhard Woeginger (TU Eindhoven). Gerhard provided me with the chapters he wrote on Polynomial Time Approximation Schemes. His material has been incredibly useful.

Without the help of a number of people from industry, it would not have been possible to produce a meaningful chapter on industrial implementations. Thanks are due to Stephan Kreipl of SAP, Shekar Krishnaswamy and Peng Qu of AMD, and Robert MacDonald of Taylor Software.

The technical production of the book would not have been possible without the invalualable help from Adam Lewenberg (Stanford University) and Achi Dosanjh (Springer). Without the continued support of the National Science Foundation this book would never have been written.

Michael Pinedo
Spring 2008
New York