Preface

Altid frejdig, når du går
veje, Gud tør kende,
selv om du til målet når
først ved verdens ende

—Chr. Richardt 1867

Since the early eighties, an active temporal database research community has sought new insight into the management of time-referenced, or temporal, data and has developed concepts, tools and techniques that better support the management of such data. Much of this activity has been motivated by the observations that most databases contain substantial amounts of temporal data and that conventional database technology offers precious little support for temporal data management.

The temporal database research community's activities are comparable to the assembly of a large jigsaw puzzle. In this research puzzle, the pieces collectively map out a variety of contributions that aim to improve the support for temporal database management. When assembling a regular jigsaw puzzle, all the pieces are expected to be available at the outset, and anybody who has tried to assemble a jigsaw puzzle knows the anguish that missing pieces can cause. A research puzzle is different. In this puzzle, it is the task of the members of the community to invent the pieces to be used in the puzzle, to shape existing pieces to make them fit with each other, and to discard pieces that turn out not to fit after all. This makes for a dynamic, ever-changing and expanding puzzle. During the past decade, pieces have been provided for the temporal-database puzzle at quite a rapidly accelerating rate, and the puzzle has now been completed to a degree where it clearly and constructively demonstrates the ability to substantially enhance the support for temporal-data management as offered by conventional technology. In fact, initial studies show that the temporal query languages proposed in this publication reduce the amount of database code needed for temporal data queries by as much as a factor of three in comparison to the standard SQL query language, and reduce the conceptual complexity by far more, thus promising a qualitative improvement in the efficiency of application development and maintenance.

This publication presents the core of my contributions to various pieces of the

temporal database puzzle, and it also marks the end of a period where almost all of my research has been related to temporal databases. This approximately eight-year period started after I completed my Ph.D. studies in January 1991, although the first contribution upon which a chapter is based was not published until some three years later. The present publication also marks the beginning of a new period where I am broadening the scope of my research to include topics such as spatio-temporal databases and general multi-dimensional databases, also termed data warehouses. Both areas are related to temporal databases, but also offer entirely new challenges and perspectives.

The contributions are divided into five parts. The first part is devoted to the semantics of temporal data, focusing most prominently on the temporal aspects of data termed valid time and transaction time, on the special notion of the current time, *now*, and on the meanings of associating values of different time data types with data. The understanding of the intricacies of temporal data as exposed in this part serves as an indispensable foundation for the remaining parts.

The next two parts concern the design of temporal query languages. The fourteen chapters in the first of these comprise a single, comprehensive case study in the design of an extension to the Structured Query Language, SQL, which offers built-in support for managing temporal data. Based on an evaluation of the resulting temporal SQL and new requirements to temporally extended query languages, the second part proceeds to propose a temporal extension to the upcoming, new SQL standard, as well as to propose an improved, "academic" temporal SQL with a design unconstrained by the standardization process.

Separate in nature from the two previous parts on query languages, but also building on the foundation laid out by the first part, the fourth part considers the design of databases containing time-referenced data. In the context of what is termed logical design, it generalizes the rich conventional relational normalization theory to apply to temporal data models and then extends this theory by offering design guidelines that take into account temporal properties of data. Covering also conceptual design, it evaluates existing temporal Entity-Relationship models, proposes a new temporal ER model, and offers a mapping of temporal ER diagrams to relational database structures, which are supported by commercial database products.

The last part is devoted to the implementation of temporal query language functionality. It builds not just on the second and third parts, but also on the first. In fact, a number of the challenges addressed here would remain undetected without an intimate understanding of the semantics of temporal data. The initial chapters focus on offering efficient support for quite specific operations on temporal databases. The later chapters emphasize the development of techniques that enable cost-effective implementation of a temporal query language in its entirety.

It is my hope that the reader will not only observe the diversity of the pieces for the temporal database jigsaw puzzle in this body of work, but also will recognize

the common themes that make the pieces interlock internally as well as with related pieces, crafted by other researchers.

The research reported here is the result of intense collaborations with a diverse group of nineteen colleagues that all share a dedication to research in temporal databases; they are listed with their current affiliations at the end of the publication, and each chapter lists its authors. Eight coauthors are in industry, and ten are with universities; ten were students during the collaborations, and three others were senior to me; eleven were in Europe, including nine in Denmark, and eight were in the United States; they cover seven nationalities; and four are female.

The interaction with this group of colleagues represents one of the privileges of conducting research and has been a source of inspiration, insight, and encouragement throughout the last eight years. Additionally, the research collaboration has had a positive effect on the relevance and general quality of the results, and I sincerely thank each of my coauthors for their parts in our joint research.

Richard T. Snodgrass has had more impact on the research reported here than any other collaborator. The research was conducted during the time since I completed my Ph.D. studies. Following the Ph.D. defense in January 1991, I packed up my car and drove from College Park, Maryland, to Tucson, Arizona, to spend the first semester of my assistant professorship at Aalborg University with Rick; in the Fall, I would return to Aalborg for a heavy teaching load. This procedure would repeat itself the following year.

From the start, Rick and I had common research interests and were immediately able to follow each other's technical reasoning with remarkable ease. Our collaboration continues to thrive to this day. As an impressionable young researcher at first and now as a more experienced researcher, I have learned and continue to learn from Rick. Being the most energetic researcher I know and a dedicated scientist, Rick never lets go of a paper before he thinks it is exactly right, which often occurs a revision or two after his coauthors are satisfied. The amount of effort needed to reach a goal is often not a real concern to Rick. He originates and takes leadership on tasks, such as the TSQL2 initiative and the SQL standardization effort, the contemplation of which alone are enough to instantly make me feel very tired. I would like to thank Rick for his enjoyable company during the assembly of the pieces of the temporal database puzzle reported in this publication, most prominently in the chapters in Parts I, II, and III. Rick has been an excellent host, both in the department and in his home, where I have stayed many a time during week-long visits. Thanks also to Rick's family: to Merrie for putting up with me, and to the kids, Eric and Melanie, who have grown up during these years.

The summer of 1994 saw my wife and me travel to Tucson, Arizona, where we were to spend my third sabbatical. This was where I first met Michael H. Böhlen. Mike had just completed his Ph.D. at ETH Zurich and had a year as a postdoc in Tucson ahead of him. I remember going with my wife LeeAnn to pick up Mike in

the airport. We were all standing at the baggage carousel when a bag the size of an American carry-on came by and Mike grabbed it. As we continued the conversation and waited, I thought, for quite a while for Mike's other luggage to appear, I commented on the slowness of the baggage delivery; at this time Mike, much to my surprise informed us that he had no more luggage for his one-year stay in Tucson. As it happened, he had only started packing a few hours before his hurried departure from Zurich and did not know what to bring anyway! This is the easy-going side of Mike in a nutshell. In research, Mike is a markedly different person. After his sabbatical, Mike accepted a position in Aalborg and has been my colleague since. He demonstrates superior technical judgment and does not compromise when it comes to quality, holding fast to his standards and ideals. I would like to extend my sincere thanks to Mike for his coauthorship on the six chapters in Part III of this publication. The powerful idea of temporal statement modifiers that underlies much of this work originates from him.

Most of my collaboration with Leo Mark dates back to the time before the period covered by this publication. I would like to thank Leo for accepting the not always easy job of being my Ph.D. advisor. I also appreciate very much the friendliness and support of Leo and his wife Inge over those years. Since that time, Leo has hosted two of my Ph.D. students and has also contributed to two chapters.

I was fortunate to get to know and coauthor one paper with James Clifford before he passed away. In addition to being a fine human being, Jim was a dedicated and inspiring scholar. I will always have fond memories of the two days I spent with Jim in Cambridge, U.K., in Spring 1994. Between conference sessions, at a cafe, in a restaurant, and in an art museum, I learned about Jim's views on our joint research area.

Finally, I would like to thank the following coauthors and colleagues for their contributions to creating a productive research environment and for putting up with my flaws: Rasa Bliujūtė, Lars Bækgaard, Curtis E. Dyreson, Heidi Gregersen, Torben B. Pedersen, Dieter Pfoser, Janne Skyt, Simonas Šaltenis, Giedrius Slivinskas, Michael D. Soo, Kristian Torp, and Nectaria Tryfona.

The research was conducted while I was employed by Aalborg University; during this time, I also spent more than two years with the Department of Computer Science at the University of Arizona. I thank both institutions for providing me with excellent working conditions. I would also like to acknowledge the Faculty of Engineering and Science, Aalborg University, for providing additional incentives to take on the task of preparing this publication. Finally, without the external funding provided by the Danish Natural and Technical Science Research Councils, the European Commission, the Nykredit corporation, and the Danish National Center for IT Research, the research reported here would not have been possible.

As mentors and role models, Bent B. Kristensen and Lars Mathiassen have had fundamental impact on my research, although their research areas are not data-

bases. Without their consistent help and encouragement, there would certainly not have been a database group at Aalborg University, and I might well have pursued a career in business rather than in academia. By his example, Bent demonstrated how it is possible to uphold strong personal standards in a political world.

Lars saw the need for increased database research in Denmark, about a decade before this became clear at a national level. It was he who formed the vision of a database group at Aalborg University, a group that is now a reality and that has served as the context for much of the research presented here. Lars embodies great human and political interest and insight; and he possesses the kind of unparalleled, awe-inspiring competence as a communicator and motivator that others can only dream of achieving.

Four colleagues helped with the technical aspects of putting together this publication. My colleague and friend over more than a decade, Frank Jensen, is a superior TEX and LATEX wizard who absolutely hates to cut corners when it comes to quality. Rasa, Simonas, and Giedrius not only helped format the chapters, but also put their growing knowledge of temporal databases to work and proof-read several of the chapters they did not coauthor.

Finally, I would like to thank my wife LeeAnn Iovanni and my parents and sister for their support and understanding. They are the ones I did not spend time with during the long days and nights of tinkering with the temporal database puzzle at the office.

Tucson, Arizona CHRISTIAN S. JENSEN
August 1999