Constraint Databases

Gabriel Kuper, Leonid Libkin, and Jan Paredaens

Gabriel Kuper
Bell Labs/Lucent Technologies,
600 Mountain Ave.,
Murray Hill,
NJ 07974, USA
kuper@bell-labs.com

Leonid Libkin
Bell Labs/Lucent Technologies,
600 Mountain Ave.,
Murray Hill,
NJ 07974, USA
libkin@bell-labs.com

Jan Paredaens
University of Antwerp (UIA),
Dept. of Math. and CS,
Universiteitsplein 1,
B-2610 Wilrijk-Antwerpen,
Belgium
pareda@uia.ua.ac.be

Preface

This book is the first comprehensive survey of the field of constraint databases. Constraint databases are a fairly new and active area of database research. The key idea is that constraints, such as linear or polynomial equations, are used to represent large, or even infinite, sets in a compact way. The ability to deal with infinite sets makes constraint databases particularly promising as a technology for integrating spatial and temporal data with standard relational databases. Constraint databases bring techniques from a variety of fields, such as logic and model theory, algebraic and computational geometry, as well as symbolic computation, to the design and analysis of data models and query languages.

The book is a collaborative effort involving many authors who have contributed chapters on their fields of expertise. Despite this, the book is designed to be read as a whole, as opposed to a collection of individual surveys. In particular, the terminology and the style of presentation have been standardized, and there are multiple cross-references between the chapters.

The idea of constraint databases goes back to the late Paris Kanellakis. While visiting the IBM T. J. Watson Research Center in 1988, he was shown a demonstration of the CLP(R) system. This system is an extension of logic programming in which rules may contain constraints over some domain, such as real or rational numbers with various arithmetic operations. Paris immediately wondered whether a database theory could be developed for such systems, analogous to the way deductive databases were inspired by logic programming. The direct consequence of this idea was his collaborative work on Constraint Query Languages [KKR90], which started the field of constraint databases. Over the years, Paris proposed many research directions that helped shape the field.

Many people helped us while we were working on the book. Erich Grädel used an early draft for a course he taught at RWTH Aachen, and many people from his Group on Mathematical Foundations of Computer Science provided useful feedback. We particularly want to thank Stephan Kreutzer, Martin Otto, Eric Rosen, and Marc Spielmann. All chapters were refereed; we thank Serge Abiteboul, Foto Afrati, Catriel Beeri, Stavros Cosmadakis, Guozhu Dong, Martin Grohe, Oliver Günther, Manolis Koubarakis, Michael Maher, Frank Neven, Kyuosek Shim, Dan Suciu, Moshe Vardi, Jeffrey Vitter, and

Limsoon Wong for their comments. Several chapters were refereed by book contributors.

We want to thank Bill Coughran, Narain Gehani, and Mihalis Yannakakis of Bell Labs and Serge Abiteboul and Sophie Cluet of INRIA for their support of the project. We thank all members of the Database Systems and Computing Principles Research departments at Bell Labs, the Advanced Database Research and Modeling group at the University of Antwerp, and the Verso group at INRIA for providing a stimulating environment and technical support. We also would like to thank Bart Kuijpers for handling the bibliography, and Hans Wössner from Springer-Verlag for editorial assistance.

Murray Hill, NJ, USA	Gabriel Kuper
Murray Hill, NJ, USA and Paris, France	Leonid Libkin
Antwerp, Belgium	Jan Paredaens
January 2000	

Table of Contents

Part II. Spatial and Temporal Data