Preface

Alberto O. Mendelzon
Computer Systems Research Institute
University of Toronto
Toronto, Canada, M5S 1A1
mendel@db.toronto.edu

This report is a collection of recent papers by members of the Hy+/GraphLog project. *Hy+* is a system for visualizing and querying databases represented by hypergraph-like structures called *hygraphs*. *GraphLog* is the visual query language supported by Hy+.

The first paper is an overview of the system. The next two papers describe applications of Hy+ in the area of distributed systems. The first one, by Consens, Hasan, and Mendelzon, is about the use of Hy+ visualizations and GraphLog queries to manipulate event traces obtained from the execution of distributed programs, particularly in debugging. In the second one, Consens and Hasan describe an approach to the visualization of computer network structure and behaviour for network management. The paper by Consens, Mendelzon and Vista explains how the deductive database systems LDL and CORAL are used to implement the visual queries provided by GraphLog. The first paper by Eigler describes the current prototype from the user's point of view, and the second one is a specification of GXF, the format used by Hy+ to store and exchange hygraphs. The three papers by Noik deal with the problem of laying out hygraphs on a computer screen (or paper): the first one is a description of a set of filters that manipulate and draw GXF hygraphs; the second one describes a new approach to the construction of "fish eye views" that give visual emphasis to selected objects in a visualization; and the third one shows how to apply this approach to hypertext documents. Finally, the paper by Mendelzon and Mendioroz presents an algorithm for storing large graphs in secondary storage so that nodes that are likely to be accessed together tend to be placed on the same disk page.

We are grateful to the Natural Sciences and Engineering Research Council of Canada, the Information Technology Research Centre, the Institute for Robotics and Intelligent Systems, and the IBM Canada Centre for Advanced Studies for their support of the work reported here.