Preface

In the two previous issues, we concentrated on papers concerning data base and data communications—exploring several systems in detail, evaluating design aids, and surveying the topics' history and current literature. This issue continues our look at data bases, emphasizing the usability and cost aspects.

Of considerable importance in any data base system is its interface to the user. The first paper, by M. M. Zloof, discusses a high-level data base language that provides the user with an effective means to define, query, update and control a data base. Mr. Zloof's paper describes Query-by-Example, a language that is currently being used experimentally. Defining his query through skeletonal tables that relate to actual data base tables, the user creates a transaction against the system easily and naturally. The paper illustrates, with many examples, the flexibility available to the user.

As data base systems develop, constraints occur which limit the ease in which the systems are used. Heyne and Daniel discuss a series of techniques which remove some of these constraints in IMS and make it simpler for the user to develop IMS applications. These extensions to IMS form a system that is a prototype of the IMS Application Development Facility. The system works as an interface between IMS and both the application programmer and user. The paper also remarks on the system's implementation.

As the use of data base systems grows, the element of cost effectiveness becomes more important. Using a relational data base, our third paper analyzes methods of evaluating a general query on the basis of the cost of accessing secondary storage. While physical storage of data and access paths are not the concern of the user, there are implications on the cost and efficiency of the system which cannot be overlooked. The data base system itself determines the efficiency of different query techniques.

While data base systems have accelerated the growth of the amount of data storage needed, other computer techniques have changed the way this storage is used. Interactive computing has led to highly dynamic data retrieval, usage, and storage. Multiple users each with their own data requirements compete for the storage resources. This has helped to create the need for automated storage management systems. Such a system is discussed in the paper on MARC: an MVS archival storage and recovery program. This paper involves a hierarchy of storage using the Mass Storage System.

Looking from the system area to that of an application area, our final paper addresses an increasingly popular computer application—econometric modeling. Models are being used to predict the industrial implications of changes in the overall national economy. This information has allowed for better forecasting of business trends. Mr. Sarma describes an IBM input-output econometric model used by the Economic Research Department at IBM. This model has allowed the study of interactions between various industries.

This is my first issue as editor of the Systems Journal. I commend my predecessor, George McQuilken, on the leadership and guidance he has given the Journal over the past three years. I look forward to working with the editorial and production staffs, as well as the many data processing professionals we rely on for assistance, so that we may continue developing the Journal as a source of new ideas in computing both from the viewpoint of systems people and application users. As the Journal is contributory, I encourage both authors and readers to participate in making the Journal an effective technical communication resource for the computing community.

Connie A. Thiel
Editor