

Preface

Herein we continue the theme of data base and data communications begun in the previous issue, with three papers, each illustrating some aspect of concepts, implementation, or design.

Much of the recent literature on data base concepts is critically reviewed by Senko. Beginning with a brief history of data structures, this paper builds to a review of recent research on such topics as standardization, performance, logical design, and future developments. For the reader wishing to study recent trends in greater detail, this survey refers to a large number of published papers.

The CICS/VS control program was developed to provide a simple programming interface for on-line terminal applications and to relieve the application programmer of the complex task of controlling terminals and communication lines. As described by Eade, Homan, and Jones, extending the implementation of CICS/VS to support Systems Network Architecture offers three advantages: existing functions are performed more effectively, new functions are provided, and a base is established for further evolution.

Designing an integrated data base may prove to be a costly and time-consuming activity. One way to reduce both time and expense is to automate parts of the design process. Raver and Hubbard describe a technique for reducing the labor of logical data base design. For DL/I structures, the technique has been implemented in the Data Base Design Aid, a program that can be used either to suggest a logical design for a data base or as a quality control tool for the experienced designer.

This issue marks my twelfth and last as editor of the Systems Journal. Over the past three years the Journal has changed gradually, I think for the better. Physically, it looks different: the covers are more colorful and more expressive; inside, more illustrations are placed in the wide margins. Content has evolved as we have added new departments: a *Preface* to introduce each issue and recognize special contributions; *Readings* for

unrefereed tutorial material; a *Forum* where technical points may be discussed; and reviews to inform our readers of significant new books.

Reader response to the Journal has remained very good. Circulation has grown steadily to over 55,000. Reprints of Systems Journal papers, not generally available three years ago, are now being ordered at a rate of 30,000 a year. A recent readership survey, conducted by an outside firm, shows that almost all recipients of the Journal read at least part of it, that half route their copies to others, and that eighty-five percent save copies for future reference. As for the Journal's value, sixty-one percent of our readers describe it as helpful in their immediate job; seventy-five percent feel it keeps them abreast of new developments in computing; and ninety percent rate the Journal excellent or good compared with other computer journals.

Producing a technical journal requires a great deal of work, not only on the part of the editorial staff, but also on the part of professionals on whom the staff relies for technical advice. Papers must be acquired and refereed. It is for refereeing—the peer review process—that the editors of the Systems Journal rely on many data processing professionals, both inside and outside IBM, who volunteer their time and energy to help produce a better journal. By tradition, the refereeing process is anonymous, so those who review, evaluate, and help strengthen each manuscript go unrecognized. As a way of acknowledging their contribution, therefore, I have listed our recent referees on page 318. To them I give my thanks, along with thanks to the editorial and production staffs of this journal, and to the many authors who have freely contributed their fine papers. I hope they all will be as helpful to my successor, Connie Thiel.

George C. McQuilken
Editor