

Preface

Interest in Open Systems Interconnection (OSI) for networks has grown with the completion of major components of these standards. As initial OSI networks are implemented, interconnection with the more than 10 000 existing SNA networks will become an increasingly common goal. The paper by Sy and his colleagues reports on a joint study between IBM and Nippon Telegraph and Telephone Corporation on OSI-SNA interconnection. It describes a possible application of OSI and compares the functions of OSI and SNA. It also introduces mapping techniques and a gateway structure, as well as discussing address translation and exception handling.

The design objectives for SNA and OSI were quite different. SNA provides a total system architecture with strength in the areas of performance and network management. OSI provides for peer-to-peer communications between self-sustaining systems that may be of different architectures. Both can be seen as complementary, and the research described in the Sy paper would support the development of future gateways between these architectures. (See also "Open Systems Interconnection" by J. R. Aschenbrenner in Volume 25, Numbers 3/4 of the *IBM Systems Journal*.)

The improved speed and storage capacity of personal workstations as well as work group systems will enable tasks to be undertaken that were formerly beyond the reach of most users. Relatively large amounts of data can be stored and processed on such systems, but it often remains difficult to discern and interpret relevant information within a number of engineering, scientific, and medical applications. By utilizing the graphics capabilities of the workstation, it becomes possible to effectively gain insight and understanding of the application data using a combination of graphical methods, windows, and animation. The paper by Farrell presents a set of imaging functions and display modes that were developed to interpret data visually for a wide range of applications.

A majority of workstation applications interface with their users by means of full-screen panels that present information and provide formatted fields for entering data. There are at the present time many vehicles available for composing these panels, and most often the sequence of panels is incorporated in the design of the application program.

The paper by Halpern, Roberts, and Lopez presents an approach which separates the definition of the presentation sequence of the panels from the application program. By the use of an *incidence matrix* the panels are treated as a graph of a network of linked nodes. This approach was implemented in a research prototype called TRYLON that enables the sequence of panel presentation to be conveniently determined outside of the application program. The result shows promise in reducing programming effort and development time.

Providing low subsecond interactive response times has proved to be extremely conducive to higher productivity on the part of system users. This was the goal of the work described in the paper by Tetzlaff and his colleagues in developing changes to the paging subsystem for VM/SP with the High Performance Option (HPO).

A primary observation was that the working set of pages used in successive interactive transactions was largely repeatable. Their paper describes the operation of the VM paging subsystem, experiments made to effect improvements in its operation, and the resulting enhancements that were incorporated in Release 3.4 of HPO. These changes provided users shortened interactive response time with an added benefit of reduced CPU time requirements.

Gary Gershon
Editor