Preface

S/390* has evolved dramatically and fundamentally from wide-ranging and powerful mainframes operating singly or in combination, to a far more wide-ranging and powerful cluster of massively parallel processors designed to handle networks and client/server systems, as well as the earlier mainframe systems. This evolution to the S/390 Parallel Sysplex* cluster has made it practical for applications and whole businesses to migrate gracefully from stand-alone environments to network and client/server environments.

The issue contains nine papers on broad architectural aspects, systems considerations, enhanced systems products, and application features of the new S/390 systems and environments. We are indebted to N. S. Bowen of the Thomas J. Watson Research Center, IBM Research Division, in Hawthorne, New York, for his considerable efforts in the creation and development of this issue. In addition, we are indebted to J.-Y. Chung, also of the Thomas J. Watson Research Center, for his noteworthy role in the early stages of issue development.

The first paper, by Nick et al., sets the stage by presenting the architectural considerations and resulting architecture for the S/390 Parallel Sysplex cluster and the OS/390* operating system. S/390 gives the new architecture its power and reach; OS/390 provides the network and client/server infrastructure for both existing and new products and applications, application programmers, and users. Key architectural aspects include advanced network data sharing, data integrity during concurrent data access, advanced distributed task execution, dynamic workload balancing, continuous availability, high performance, and increased cluster and network efficiencies.

An essential feature of highly distributed, network-based, and clustered systems is their ability to share data effectively, efficiently, synchronously, with high throughput, but without loss of integrity. Bowen et al. discuss the S/390 and OS/390 advanced multisystem locking architecture and models that allow a new, deeper level of data sharing, while specifically supporting both update-intensive and read-intensive applications.

The long-term value of distributed, networked, and clustered architectures has become dependent, in a fundamental way, on their ability to scale up and down based on business and application requirements. A key technical feature of scalability is the linearity or near linearity of transaction rates as nodes are added, despite the additional systems management burden. Many of the S/390 architectural capabilities for scalability are found in the coupling facility (CF). The authors—King, Dias, and Yu—describe and quantify the CF and related S/390 cluster architectural features that provide system-wide scalability.

Managing and balancing a distributed system and its workload is a reflection of both hardware and software capabilities. Aman et al. focus on a crucial aspect of the software side: adaptive algorithms that perform workload balancing and reflect the users' business requirements. The authors provide a complete discussion of the OS/390 adaptive algorithms and their genesis, along with two extended examples of their business-oriented adaptive power.

Networks and clusters of computers have become and will become ever larger, more dispersed physically, more specialized computationally, and more critical to the success of businesses and applications. As a consequence, there is a demand for higher availability, less down time, and fewer systems requirements for human intervention in systems management. Bowen et al. treat an important aspect of this evolution: automatic process restart. The authors discuss the technical requirements for automatic process restart, its impact on systems design, and the features in and on OS/390 that support this function.

System management requirements and their realization in the CICS* (Customer Information Control System) and the CICSplex in the S/390 context are the subjects of a paper on the CICSPlex* System Manager (SM) and related products, written by Johnson. One key to the advanced system management features provided by SM is the creation of a single-system image that makes resource locations transparent for the purposes of high-level systems control.

Josten et al. present the relationship between the relational database management power of DB2* (DATABASE 2*) for OS/390 and the data-sharing power of the S/390 CF. DB2 has been enhanced and extended to provide database management in the networked and clustered environment made possible by S/390 and OS/390. A new DB2 data-sharing function was developed to take advantage of the CF, while providing advanced network features for DB2 users.

Transaction processing across an S/390 Parallel Sysplex cluster presents the opportunity to provide enhanced operations to new and existing applications. Banks, Davies, and Moxey discuss how that opportunity has been seized in the case of CICS/ESA*, resulting in a richer transaction-processing environment with reduced complexity through a single-system view, while providing enhanced capabilities transparently to existing applications.

Strickland explores the extension of and enhancements to VSAM (Virtual Storage Access Method) for the new S/390 environment. The author's focus is on the ways in which VSAM can utilize the data-sharing capabilities of S/390 CF to provide serialization for VSAM data streams. As with many other system services that have been and are being made available for S/390 Parallel Sysplex clusters, the VSAM architecture, design, and technology had to be revisited, understood in the new context, and ported or modified for its enhanced role in S/390.

The next issue of the ***Journal*** will present papers on a variety of subjects, including use of S/390 for multimedia, on-line reorganization of databases, and new uses of networked computers.

Gene F. Hoffnagle
Editor