Reliable Distributed System Software

John A. Stankovic

IEEE CATALOG NUMBER EHO230-3
LIBRARY OF CONGRESS NUMBER 85-60382
IEEE COMPUTER SOCIETY ORDER NUMBER 570
ISBN 0-8186-0570-7

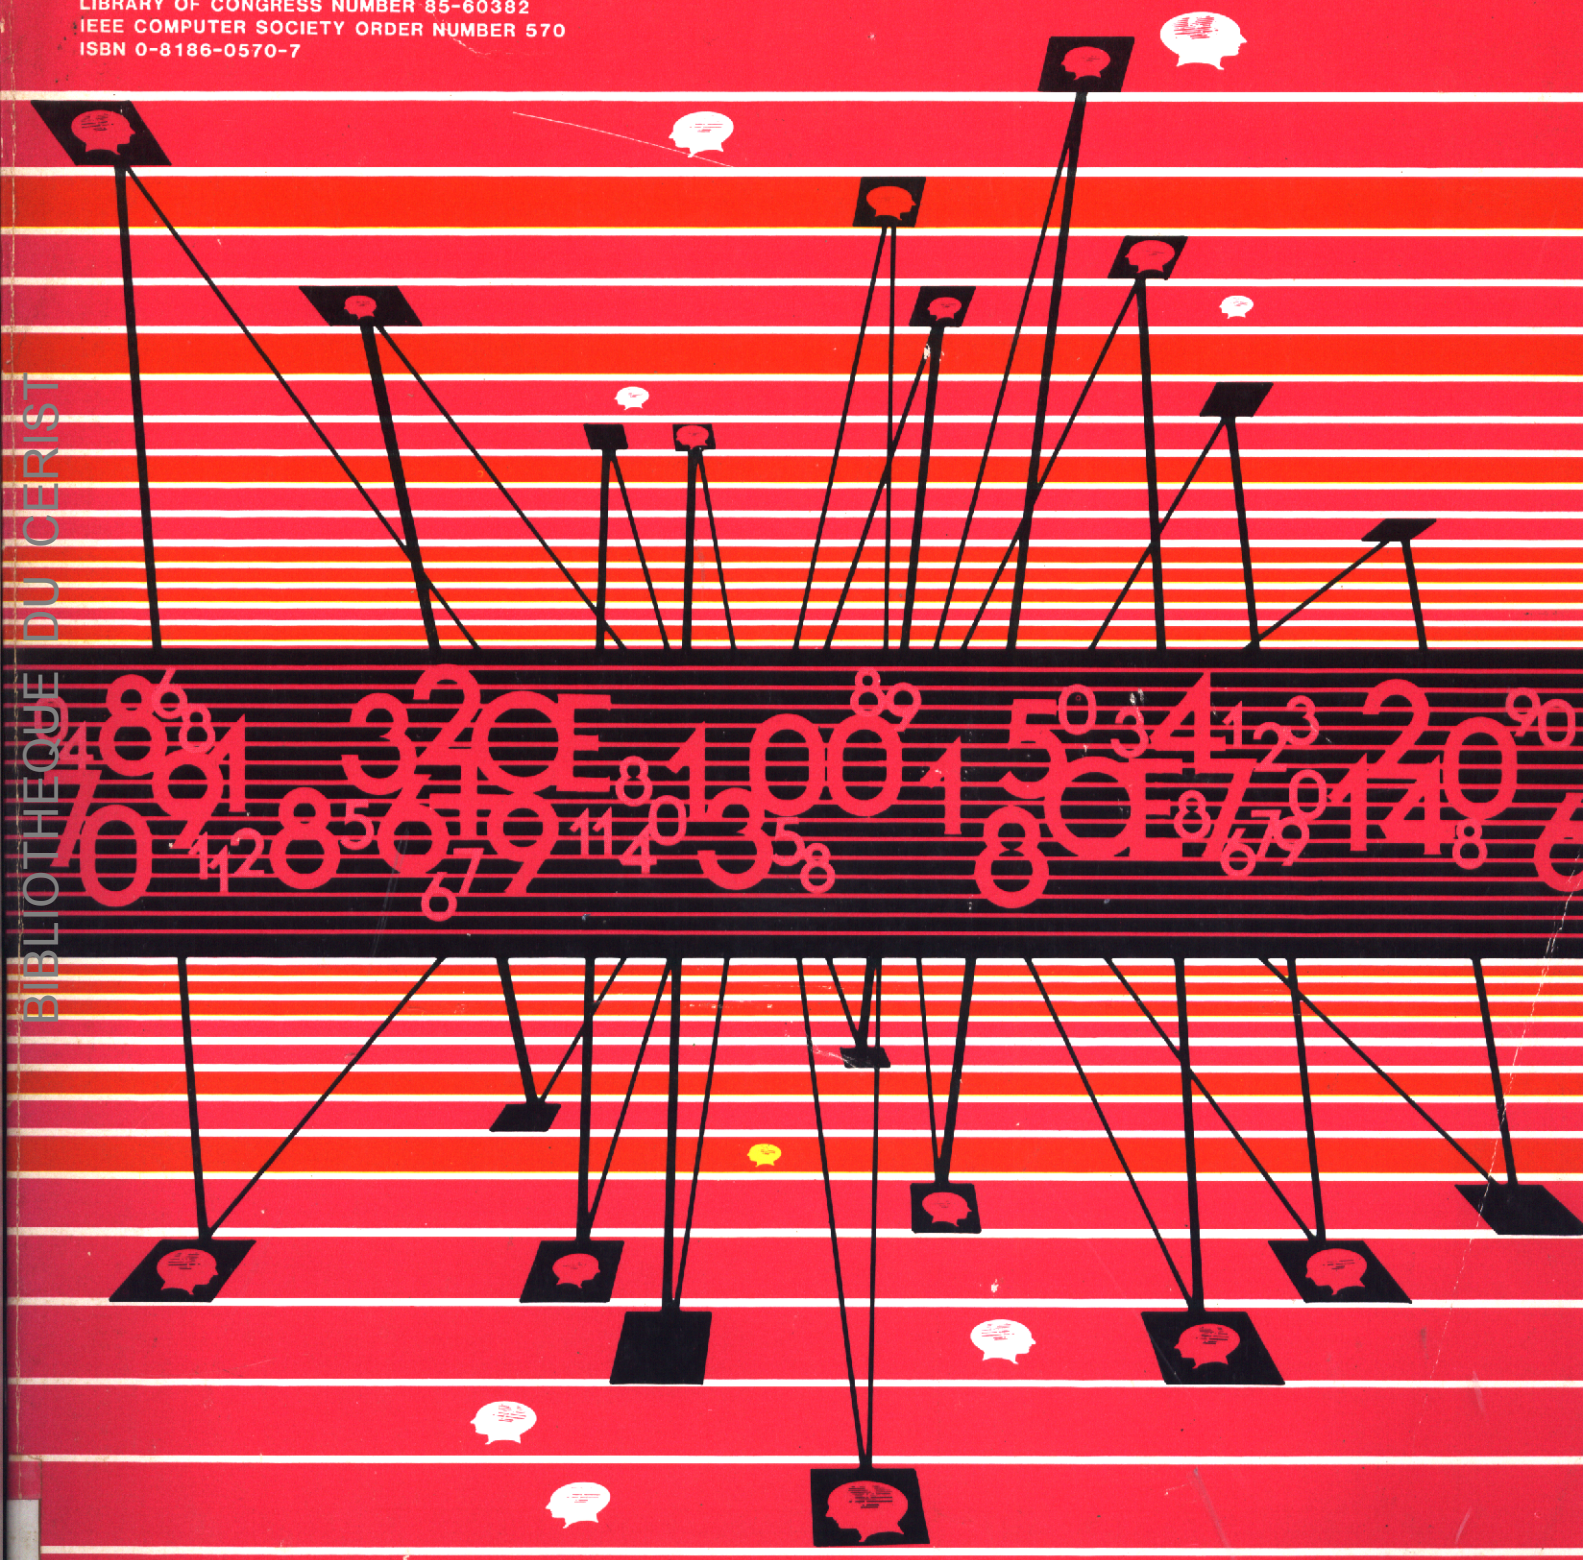

IEEE COMPUTER SOCIETY
IEEE COMPUTER SOCIETY PRESS
IEEE
THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Reliable Distributed System Software

John A. Stankovic

IEEE CATALOG NUMBER EH0230-3
LIBRARY OF CONGRESS NUMBER 85-60382
IEEE COMPUTER SOCIETY ORDER NUMBER 570
ISBN 0-8186-0570-7

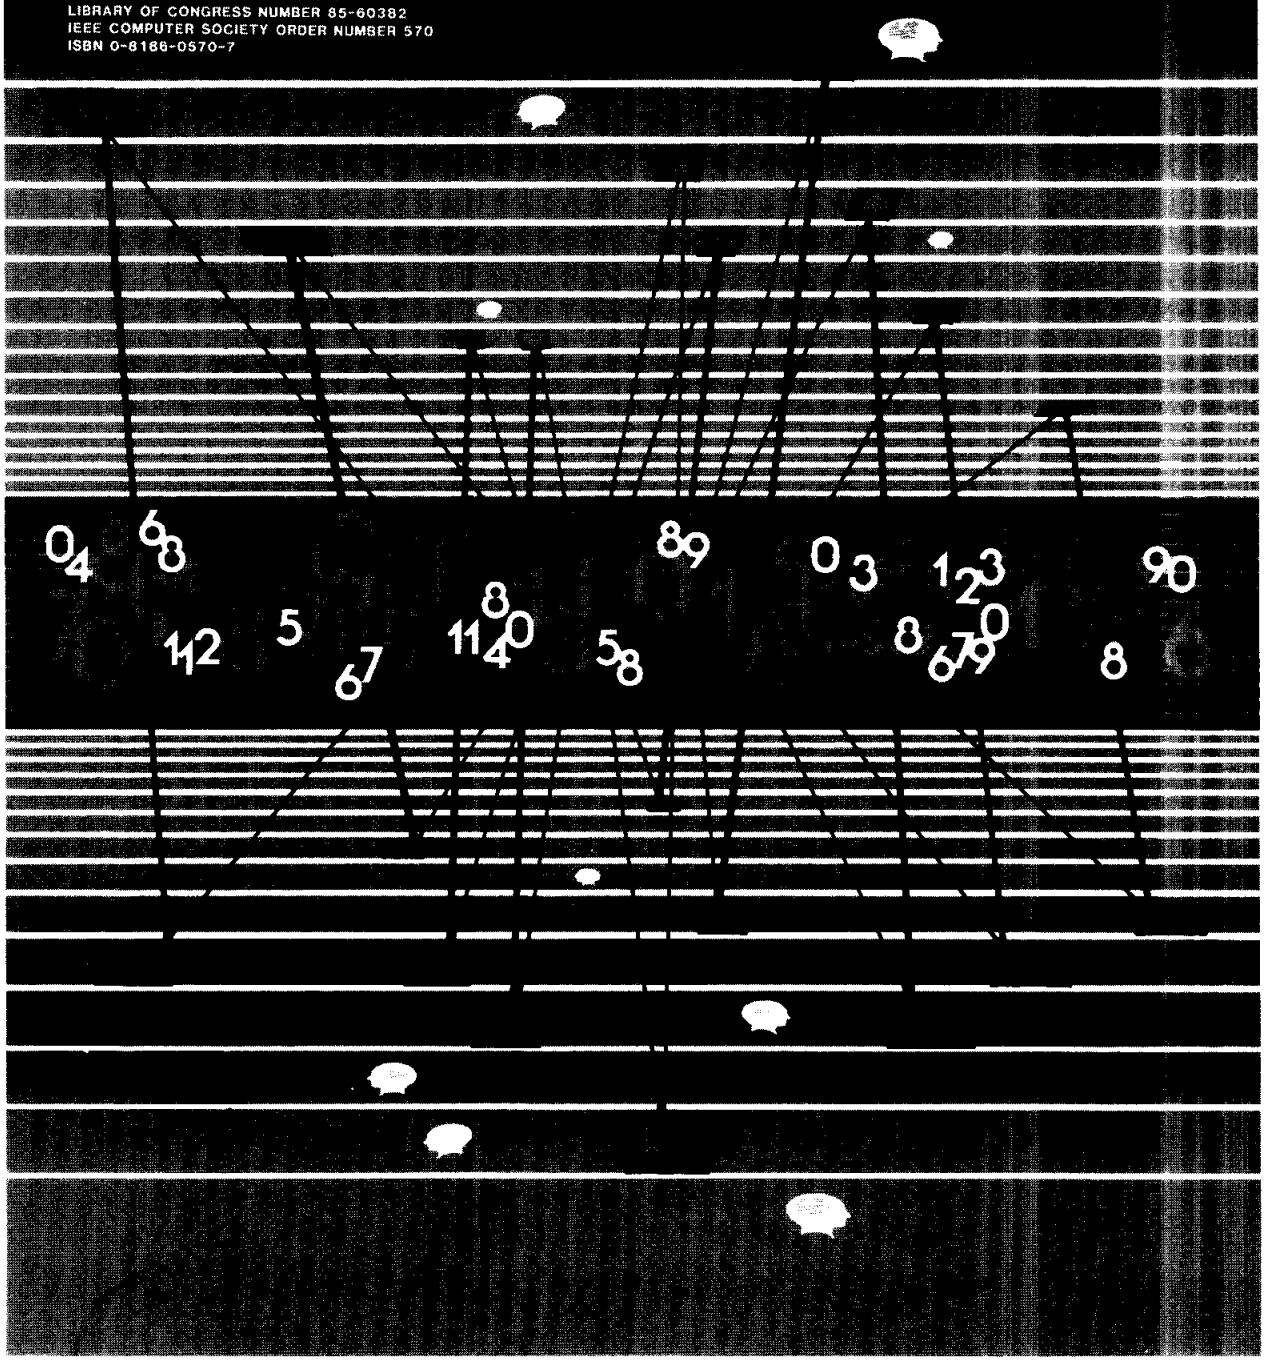

IEEE COMPUTER SOCIETY
IEEE COMPUTER SOCIETY PRESS
IEEE
THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Published by IEEE Computer Society Press
1109 Spring Street
Suite 300
Silver Spring, MD 20910

IEEE Catalog Number EHO230-3
Library of Congress Number 85-60382
IEEE Computer Society Order Number 570
ISBN 0-8186-0570-7 (Paper)
ISBN 0-8186-4570-9 (Microfiche)

Order from: IEEE Computer Society
Post Office Box 80452
Worldway Postal Center
Los Angeles, CA 90080

IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08854

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Acknowledgements

Some of the material found in Sections 1 and 2 was jointly developed by Prof. Kohler, Prof. Ramamritham, and me. I thank them for their work in these areas.

Preface

This tutorial text organizes and presents issues and concepts related to reliability in distributed systems software. In this presentation, software includes communication protocols, logical interprocess communication (IPC) support, distributed programming languages, distributed operating systems, and distributed databases. While there exists a considerable body of knowledge on reliability techniques for hardware, communication protocols, and distributed databases, there has been significantly less work on reliability in distributed operating systems. In fact, much of the distributed operating system work only briefly or implicitly addresses reliability. This tutorial text attempts to cover the broad spectrum of reliability techniques used in distributed system software including distributed operating systems.

A reader can expect to learn what reliability is, what reliability techniques are used in the different areas of distributed system software, and how reliability techniques can be better applied across all areas of distributed systems software (especially in the distributed operating system area). The text is primarily written for computer scientists and systems programmers who need to understand, design, or implement reliable system software. Hardware designers also can benefit from knowledge of the software reliability techniques by obtaining a better overall system perspective on reliability.

The text first presents an overview of reliability (hardware and software) as well as an overview of "general" distributed computer systems research. Reliability techniques for each subarea of distributed systems software is then presented including: the communication subnet (Section 3), the operating system area (Sections 4–7 inclusive), and the database area (Section 8). Techniques in the operating system area are divided into subareas including logical IPC and distributed programming languages (Section 4), distributed control (Section 5), structuring distributed systems for reliability (Section 6), and a summary collection of software reliability techniques (Section 7). The text concludes with several case studies of reliable systems (Section 9). An extensive set of references and a bibliography are also provided.

John A. Stankovic

Table of Contents