AN INTRODUCTION TO DISTRIBUTED SYSTEMS

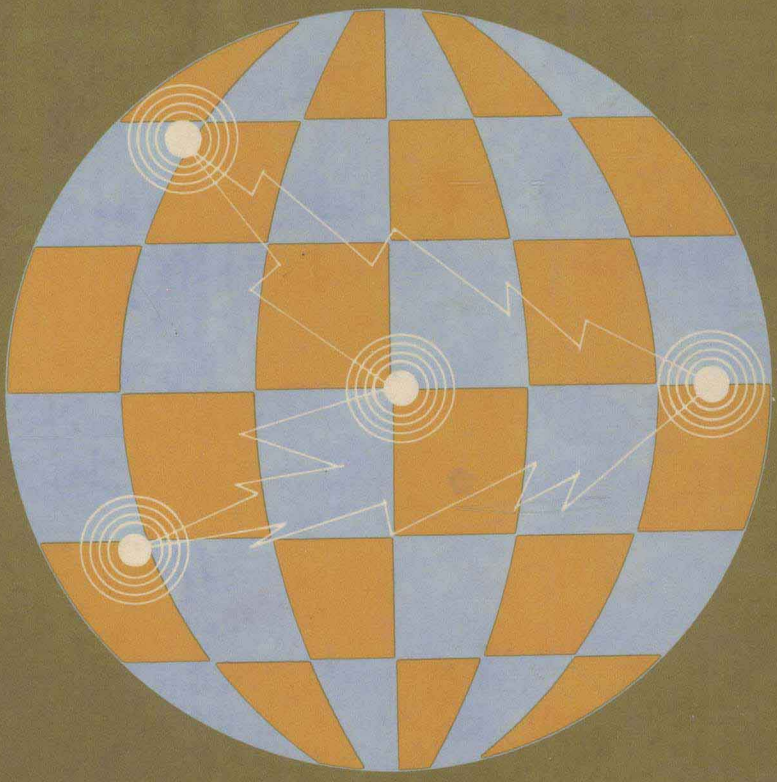

Bennet P. Lientz

AN INTRODUCTION TO DISTRIBUTED SYSTEMS

Bennet P. Lientz
Graduate School of Management
University of California, Los Angeles

ADDISON-WESLEY PUBLISHING COMPANY
Reading, Massachusetts • Menlo Park, California
London • Amsterdam • Don Mills, Ontario • Sydney

Reproduced by Addison-Wesley from camera-ready copy supplied by the author.

Library of Congress Cataloging in Publication Data

Lientz, Bennet P.
An introduction to distributed systems.

Includes index.
1. Electronic data processing--Distributed processing.
I. Title.
QA76.9.D5L52 001.64 80-23452
ISBN 0-201-04297-5

Copyright © 1981 by Addison-Wesley Publishing Company, Inc. Phillipines copyright 1981 by Addison-Wesley Publishing Company, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

ISBN 0-201-04297-5
ABCDEFGHIJK-DO-89876543210

Preface

Distributed systems are becoming a major focus of industrial and governmental computer efforts. A wide variety of choices is available. More analysis is required. Management and user issues are more important, requiring more analysis. Selection decisions are more complex. Users are more deeply involved. The risks of failure are higher. Control and security issues must be addressed. Overlaying this has been the rapid pace of technological advance.

All of these developments indicate that we need new approaches to computer systems analysis at the college and professional levels. We need to be able to deal with change as well as computer and communication concepts. Technology must be viewed as transient--not stable. We also must address more of the system life cycle. Specifically, we need to consider maintenance and enhancement.

This book addresses distributed systems in its broadest sense. Only a minimal computer course is assumed. Although it is not a traditional analysis book, many systems analysis methods are similar to those used for traditional central systems. The differences lie in the application of the methods and what must be considered to avoid failure. Avoiding failure and identifying potential problems are major foci of this book.

Systems analysis books are often dry and, although many times accurate and complete, fail to get the message across. Thus, we have attempted here to write the material in an easier, more informal style. We also have included case studies, examples and appendices as background. We consider a general systems arena throughout the book. As part of this arena an area that is likely to be a force for substantial future change is office systems. We explore the environment of the office and address specific analysis and design problems.

Throughout the book we not only consider techniques, but also the process. Have you ever had students come up and ask, "How is this used? When is it used?" We are going to look at both the managerial and technical environment. The system life cycle, the roles in the life cycle, and techniques combine results in recent literature as well as established methods. This is a book on how work is done in real life and the interaction between managerial and technical people. It contains a discussion of many techniques, but it is not a grab bag of tools.

The book is organized around a one quarter or one semester course at the college level. Part I introduces the basic concepts, discusses computer and communications, and describes the planning process. The appendices on terms, references and office systems can be referred to at this point in class. The remainder of the book address the techniques employed during the life cycle. Specifically, Part II addresses analysis, design and construction while Part III addresses installation, operation and maintenance. The book closes with a chapter providing overall perspective to the material, as well as two case studies (Part IV). The approach is based on acquiring systems through a Request for Proposal process as well as construction from scratch.

此为试读,需要完整PDF请访问: www.ertongbook.com

A great debt is owed to the many students at UCLA and in professional seminars as well as to several colleagues - Eph McLean and Burt Swanson - who provided suggestions and encouragement. The materials were tested and used for the last two years in classes at U.C.L.A., as well as in several techical seminars. We are also indebted to managers in several corporations who consented to having their situations used as cases. Thanks are also due to the two reviewers who made many valuable suggestions. The comments of Dr. George L. Muer were most helpful. Finally, we wish to express our thanks and gratitude to Ms. Jeanette Emerson, who perservered and tolerated this project for so long.

Contents

Part I Technology and Planning

Chapter 1 Introduction

An Overview of Distributed Systems

The computer field is frequented with buzz words and jargon. One phrase is "distributed processing." Another is "computer network." What do these mean? Are they really important? Some computer professionals have touted them as concepts which will revolutionize computing. But will they? The answer is, in most cases, no. Computer systems exist to serve organizations and business processes and systems. These systems change slowly. Supporting accounting and computing systems also change gradually. In an evolutionary - not revolutionary - way.

What are distributed systems? In a broad sense they are dispersing the handling of information to various organizations and users. The definition is

oversimplified, but it is useful as a start. Is the concept new? In this general sense, the answer is no. Let us consider the beginning of computers. Prior to computerization, most companies dispersed functions to specific departments. Much accounting was done in branch banks or local regional offices. Accounting machines and tabulators were used to support these operations. Summary data were fed to central offices for analysis and further tabulation. Communication was by telegraph, telephone or courier. In our definition, this is distributed processing.

Then came the computer. Operations were semi-centralized. Claims processing was done around a regional base, as was the clearing of checks. We had entered the computer age. It was the late 1950's and early 1960's. Economies of scale began to be realized. Three or more regional centers could be combined into one larger center for efficiency.

At this time, with the arrival of large computers on a commercial scale, we saw the emergence of centralization. After all, to feed and work with this computer you were advised to centralize staffs around the computer. The idea of a large corporate computer center serving users on a global basis was born. Many such centers were built. Many more were planned and never built.

Why did the centralizing of computers fail in some cases? Was it the emergence of minicomputers and telecommunications? In many cases it was not! Frequently, centralized facilities fail because of some or all of the following reasons:

- Large central staffs can become isolated from user organizations;
- Larger organizations for projects can become unwieldy;
- Central organizations can fail to keep up with user demands;
- Procedures are excessively formalized.

此为试读,需要完整PDF请访问: www.ertongbook.com

This is supported in many case studies. A classic example lies in Brooks' Mythical Man Month (Reading, Mass.: Addison-Wesley Publishing, 1976) where he points out managerial and organizational problems associated with the development of the large IBM S/360 operating system.

Frustrated managers, users and systems groups see in distributed processing a potential solution to increase effectiveness and performance. Distributed computing can involve moving terminals and computer power to different locations. These could be dispersed to regions or located with users. Functions can be moved out as well.

Because of the appeal of distributed processing, some feel that all functions should be distributed. This is not really the case. Central messages of this book are:

- Distributed computing offers more options for configuring systems.
- The increased options lead to more complexity--hardware, software, communications, and organization.
- Whether or not an application should be distributed and the degree of distribution is a business decision as much as, if not more than, a technical decision.

This discussion can be shown graphically. In Figure 1.1 we can see the historical trend of centralization. Figure 1.2 shows a decision-type table based on appropriate actions. The second figure is useful here, but we will see that there is a continuum of degrees of centralization--so that this figure is a simplification. As centralized facilities were at one time "the only game in town", most computer applications would be in areas A and B. Whether an application should be in A or B is a decision based on the user environment, technology, etc. The pressure toward distributed systems tends to push applications from A and B to C and D. What we want to do is to ensure that most applications move into the diagonal boxes--A and D.

Figure 1.1

Centralization

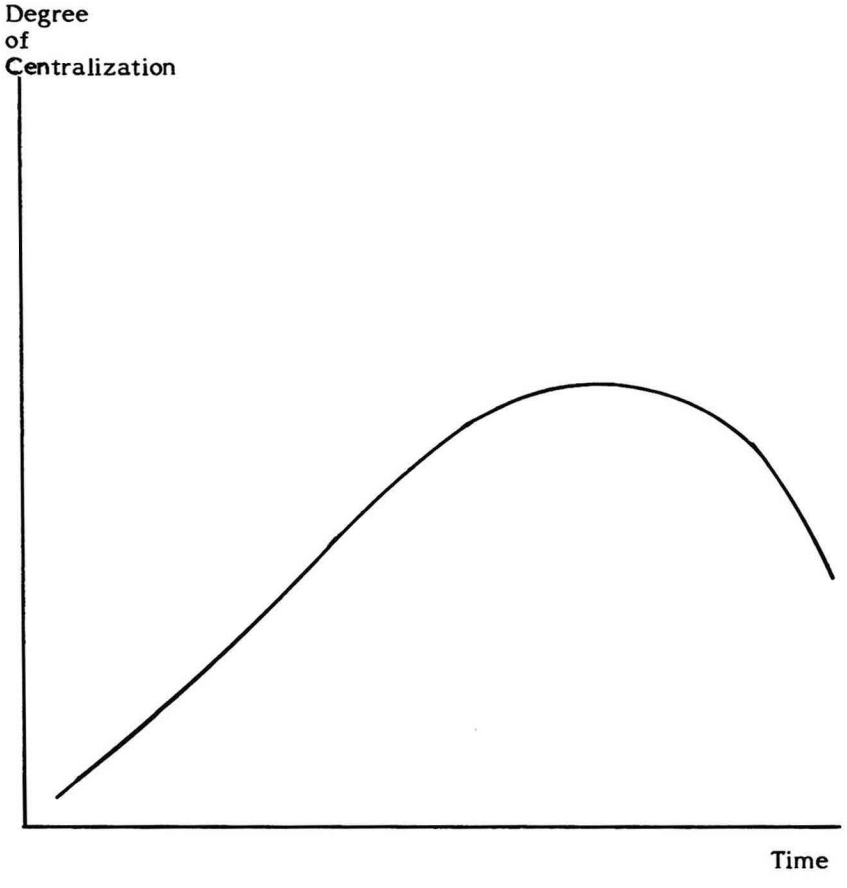

Figure 1.2

Application Classification

	Should be Centralized	Should be Distributed
Currently Centralized	A	B
Currently Distributed	C	D

Who is Involved

Traditionally many central systems have had little direct contact with user organizations and management. Time cards, purchase orders, etc., would be collected, keypunched, verified and put into the machine. The output reports would be printed and distributed. The cycle would be repeated--a standard batch processing mode. In these applications the following organizations are involved:

- o Users -- prepare input forms, read and use output reports;
- o Operations
  - Keypunch and verify data
  - Run production job
  - Perform backup and retention of data
  - Control distribution of data and reports;
- o System development
  - Maintain applications
  - Enhance application
  - Develop new systems;
- o Management -- utilize reports through intermediaries.

Figure 1.3 shows some of the information flows between these roles.

The above situation of multiple functions not only applies to large computer operations, it also applies to home microcomputers since all of the above functions are performed by one person.

Now consider a distributed system. In one situation we might have:

Figure 1.3

Examples of Relationships
Between Roles

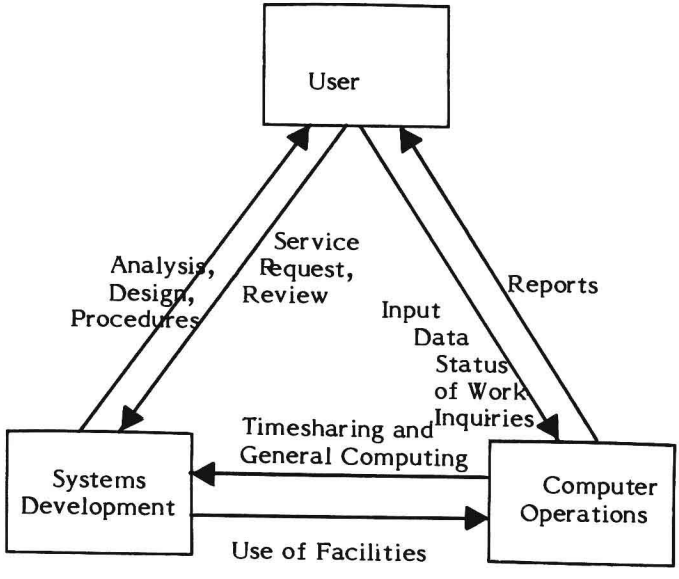

- Users

    - Prepare data
    - Input, edit data
    - Initiate processing
    - Read reports
    - Perform simple programming to retrieve data or generate simple reports;

- Operations (usually a very small staff)

    - Run production
    - Perform backup, retention of data;

- System development -- perform major maintenance and enhancement;

- Management -- utilize reports through an intermediary.

Notice the difference. For management it is the same. For the users the responsibilities have increased. Users are doing more of their own processing and simple programming.

What is the impact of the difference in roles? The arguments in favor of distributed systems point to effectiveness. The users are doing the work without intermediaries in operations. There may be a reduction in operations staff and potentially in system development since users are doing their own simple programming. The opponents of distributed computing would ask the following questions. How will data integrity and security be maintained? Will the users willingly take on the additional work? How will the effectiveness be measured? What is the organizational impact? We will explore these issues in later chapters.

It is of interest to note that there can be an almost unlimited set of options. Not all of these functions have to be distributed. The hardware does not need to be distributed.

此为试读,需要完整PDF请访问: www.ertongbook.com