

DISTRIBUTED DATA BASES

Edited by
I.W. DRAFFAN and F. POOLE
Sheffield City Polytechnic

CAMBRIDGE UNIVERSITY PRESS
CAMBRIDGE
LONDON NEW YORK NEW ROCHELLE
MELBOURNE SYDNEY

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
32 East 57th Street, New York, N.Y. 10022, USA
296 Beaconsfield Parade, Middle Park, Melbourne 3206, Australia

© Cambridge University Press 1980

First published 1980

Printed in the United States of America

Printed and bound by The Book Press, Brattleboro, Vermont

Library of Congress Cataloguing in Publication Data

Distributed data bases.

1. Data base management. 2. Electronic data processing -
Distributed processing. I. Draffan, I.W. II. Poole, F.
001.6'442 QA76.9.D3 80-40399

ISBN 0 521 23091 8

DISTRIBUTED DATA BASES

LIST OF COURSE LECTURERS

Professor G. Bracchi
Istituto di Elettrotecnica ed Elettronica
Politecnico di Milano, Italy

Mr I.W. Draffan
Department of Computer Studies, Sheffield City Polytechnic,
Pond Street
Sheffield, UK

Dr J.M. Gross
Data Logic, Westway House
320 Ruislip Road East
Greenford, Middlesex, UK

Mr P.E. Jackson
Data Logic, Westway House
320 Ruislip Road East
Greenford, Middlesex, UK

Mr J. Joyce
Data Logic, Westway House
320 Ruislip Road East
Greenford, Middlesex, UK

Professor P.J.H. King
Department of Computer Science
Birkbeck College
University of London, UK

Dr G. le Lann
IRIA, Projet SIRIUS
BP.105-78150 Le Chesnay, France

Dr B. Lindsay
IBM Research Laboratory
5600 Cottle Road, San José, California 95193, USA

Mr F.A. McGuire
Data Logic, Westway House
320 Ruislip Road East
Greenford, Middlesex, UK

Dr F. Poole
Department of Computer Studies
Sheffield City Polytechnic, Pond Street, Sheffield, UK

Professor F.A. Schreiber
Istituto di Elettrotecnica ed Elettronica
Politecnico di Milano, Italy

Dr P. Selinger
IBM Research Laboratory
5600 Cottle Road, San José, California 95193, USA

Dr S. Spaccapietra
Institut de Programmation, Université Pierre et Marie Curie (Paris 6)
4 place Jussieu, 75230 Paris Cedex 05, France

Professor M. Stonebraker
Electronics Research Laboratory, College of Engineering
University of California, Berkeley, California 94720, USA

Dr F.E. Taylor
PO Box 5, Knutsford, Cheshire, UK

PREFACE

Distributed Data Bases are the result of the convergence of apparently different technologies; distribution representing decentralisation, data bases representing centralisation and networking representing communication.

Distributed Data Bases introduce new problems and exacerbate others which are well known in non-distributed contexts. It is likely that many future data base implementations will be of the distributed type and there is therefore a need for a comprehensive review of this important subject. This text is based on an Advanced Course on Distributed Data Bases organised by the Editors and held at Sheffield City Polytechnic in 1979. The course was sponsored by the Science Research Council in the United Kingdom and by the Informatics Training Group of CREST in Europe.

This event brought together a number of distinguished speakers and research workers, from Europe, the United Kingdom and the United States of America, as well as eighty participants drawn from business, commercial, industrial and academic institutions in the countries of the European Economic Community. The course focussed on the problems associated with the set-up and administration of distributed data base systems, with the needs of organisations and with the deeper technical issues. This text develops a comprehensive framework of understanding against which the problem areas can be identified and classified, and attempts to suggest potential solutions whenever possible. The original course lecture notes have been extensively revised by the authors and editors for publication in its present form.

ACKNOWLEDGEMENTS

The Editors gratefully acknowledge the support of their wives Gwen and Christine during this project and express their particular thanks to Miss Linda Slater for her excellent secretarial support.

CONTENTS

	Page
List of Course Lecturers	vii
Preface	ix
1. Distributed systems in perspective Frank E. Taylor	1
2. A critical review of aspects of database theory P.J.H. King	33
3. The classification of distributed data base management systems I.W. Draffan and F. Poole	57
4. Distributed query processing G. Bracchi, C. Baldissera and S. Ceri	83
5. Homogeneous distributed data base systems Michael Stonebraker	103
6. Heterogeneous data base distribution Stefano Spaccapietra	155
7. Consistency, synchronisation and concurrency control G. Le Lann	195
8. Replicated data Patricia G. Selinger	223
9. Authorization and views Patricia G. Selinger	233

10. Single and multi-site recovery facilities	
Bruce G. Lindsay	247
11. Distributed data base design and administration	
J.M. Gross, P.E. Jackson, J. Joyce and F.A. McGuire	285
12. Distributed data base applications: an overview	
F.A. Schreiber, C. Baldissera and S. Ceri	323
13. Distributed data base applications: model of an access strategy in a distributed data base	
G. Pelagatti and F.A. Schreiber	337
Appendix The argument against CODASYL	
Michael Stonebraker	361
Index	371

CHAPTER 1

DISTRIBUTED SYSTEMS IN PERSPECTIVE

Frank E. Taylor

Principal, Systems Technology Consultants,
PO Box 5, Knutsford, ~~Cheshire~~, UK

1. INTRODUCTION

Distributed computing systems are a new type of system. They emerged from a successful combination of powerful small time computers and traditional mainframes. In the broadest sense, all resource-sharing systems could be described as distributed systems when two or more processors are involved. In its narrowest sense, the term refers specifically to systems which combine small and large computers [1]. From a data base point of view, the need to arrange for geographically separate systems to intercommunicate has led to an acceleration of development. In particular, increasing staff costs have led to a progressive migration of computing interfaces right into the user environment. This trend is in line with the present policies concerned with conservation of energy, especially when one compares the kilowatts required to take a commuter to his office with the milliwatts required to take data to the same user if he were to work at home or at a workstation close to his home.

As a consequence, it is perhaps necessary to say that the balance between system elements, particularly hardware and staff, is changing rapidly. Figure 1 shows qualitatively the differing balance of costs between hardware, the cost of which is dropping rapidly, and staff costs which are rising steadily as a consequence of inflation and scarcity of staff. As a consequence, it is becoming more and more desirable to make the best use of costly staff time by migrating towards a hardware intensive systems solution.

2. THE NATURE AND CLASSIFICATION OF DISTRIBUTED SYSTEMS

Distributed systems are systems which geographically distribute elements within a computing system. The two main elements which can be distributed are processing power or storage or both. The following is thus a definition of a distributed system:

"A series of separate but interworking subsystems which have processing power or storage or both".

F.E. TAYLOR

Relative trends - the cost of hardware and staff




Fig.1. The balance is changing and staff costs are increasing in importance.

Note that some but not necessarily all subsystems are usually at or near the point of user activity, and the communications between active subsystems is usually provided by on-line links, not off-line links. Many, although not all, distributed systems consist of small computing systems interworking with larger systems (the systems local to users are usually small systems, the capability of which is complemented by remote facilities) usually in the form of a large mainframe or a gateway to a computer network.

The philosophy of distributed computing has only become feasible since the early 1970s through the emergence of low-cost, yet relatively powerful, small mini and micro computers at a price which makes it possible and feasible to distribute large numbers of such computers across a given organisation. The availability of such processing power, in combination with associated low-cost data storage devices has enabled processing power and storage to be placed within the user environment close to the point of user activity.

In order to qualify for use of the term 'distributed system', small computers at or near their users must at least carry out

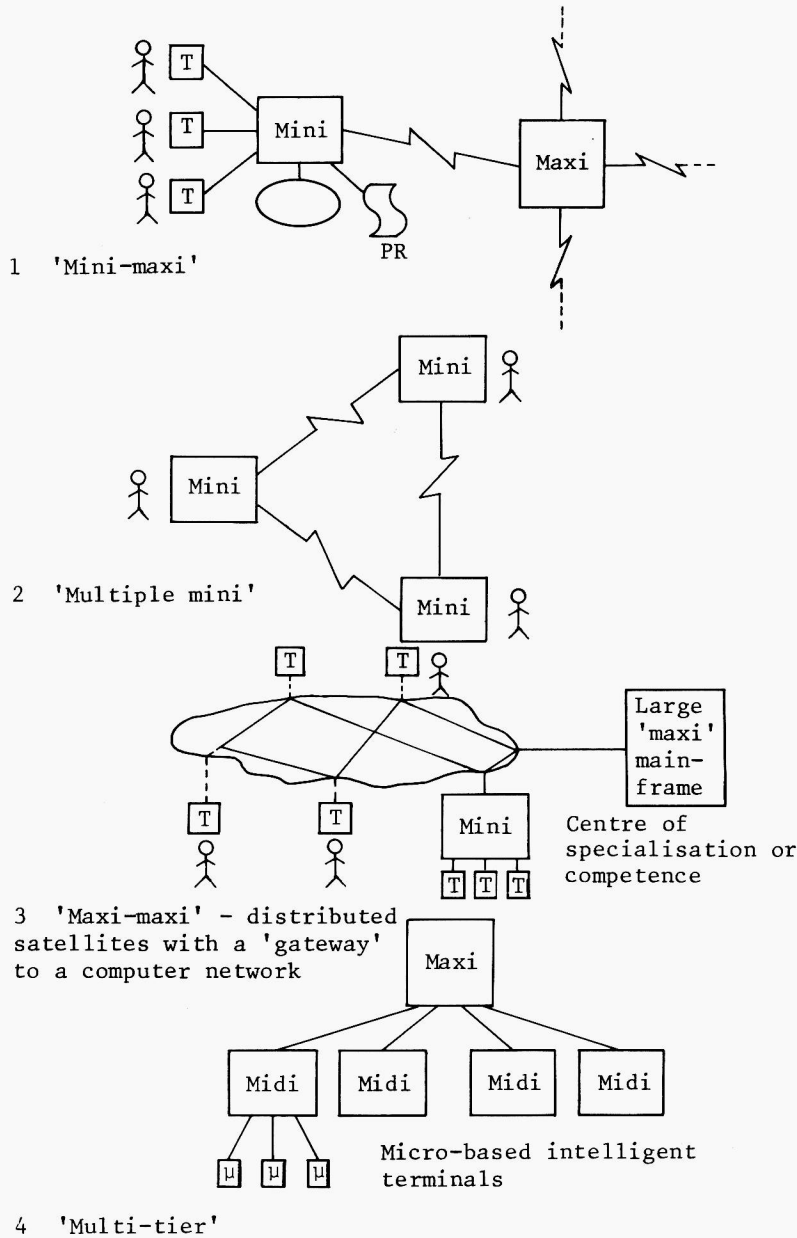


Fig.2. Some types of distributed system.

data validation functions, i.e. perform a user task at the lowest level. Those which carry out communications functions only do not qualify. Some typical types of distributed systems are shown schematically in Figure 2. (See also [2].)

As can be seen, the first type of system involves a combination of a minicomputer or number of minicomputers with a large computer. In this case, it is usual for activities within the total system to be divided in such a way that an appropriate computing element handles work for which it is particularly suited. Generally speaking, small high-volume jobs are handled locally at the point of user activity by the minicomputers, whilst more complex tasks and those which have to be centralised for administrative reasons are normally handled by the larger mainframe.

The second type of system involves interworking minicomputers. This can involve load sharing between the active elements, although this is only usually necessary under overload condition at any one particular active location. Automatic transfer of jobs to a remote location is not usual in such systems, unless the active subsystems have widely differing capability.

A third type of system shown involves interlinked resource sharing mainframes, which may subdivide the processing of user tasks between two or more such host mainframes. Whilst this is distributed computing in its broadest sense, it is much closer to computer network technology, than the technology of distributed systems in the narrower sense defined above.

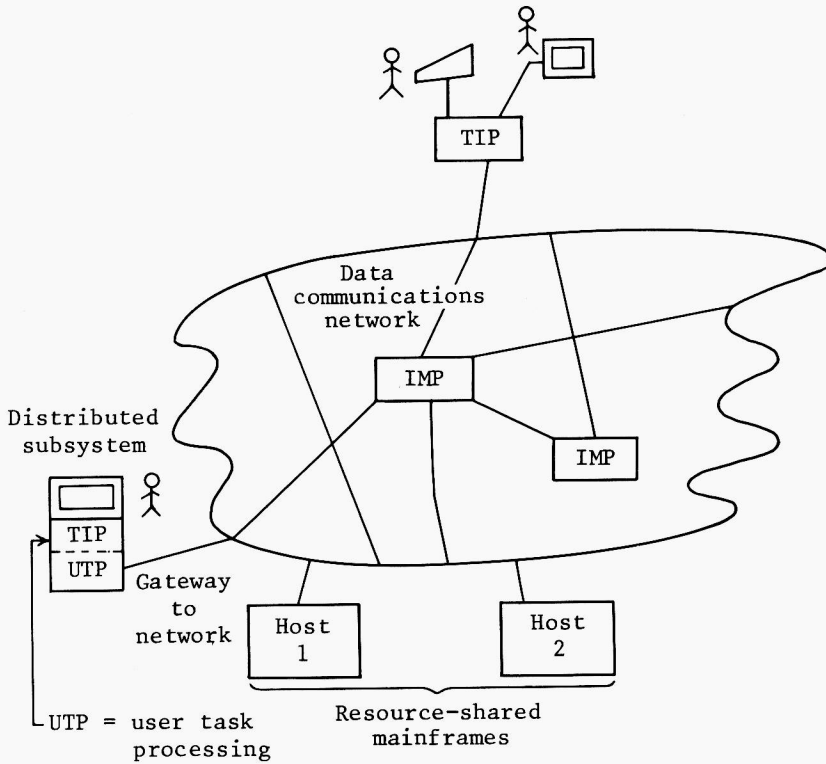
The last type of system shown, a multi-tier system, is an extension of the mini/maxi configuration just outlined. In the case shown, the minicomputer clusters have been replaced by regional 'midi' computers in combination with small local micro-computer-based terminals, which in this case are usually distributed at strategic points across the user environment.

3. THE DISTINCTION BETWEEN DISTRIBUTED SYSTEMS AND RESOURCE-SHARING COMPUTER NETWORKS

Figure 3 shows the fundamental distinction between resource-sharing computer networks which centralise processing, and distributed processing systems. In the case of distributed systems, processing power is fragmented in such a way that fragments of the processing and storage elements within a system are taken into the user environment. In contrast, within resource-sharing computer networks, the accent is on efficiently and effectively transporting tasks to a suitable host resource. In this case the terminals act as a gateway to the system, but do not normally carry out user task processing at or near the points of user activity.

DISTRIBUTED SYSTEMS IN PERSPECTIVE

The distinction between resource-sharing computer networks, which centralise processing, and distributed processing systems



Schematic diagram of a networking system (1970 on). Accent is on efficiently transporting jobs to a suitable 'host' resource.

In contrast, distributed systems take all or part of resource to job.

* Note that a network can have a distributed element if terminals process minimal tasks, or processing shared between hosts.

Fig.3. Resource-sharing systems contrasted with distributed systems

The distinction becomes blurred in cases where small distributed computing subsystems link to a computer network via an appropriate gateway, so that they can off-load tasks which they either cannot process or which would create an overload condition, to an appropriate host resource via a data communications network.

4. SOME ADVANTAGES OF DISTRIBUTED SYSTEMS

4.1 Aligning with geographically distributed operations

Many organisations have geographically dispersed operations. The distributed systems philosophy allows separate satellite subsystems to be placed within geographically separate locations, yet linked in such a way that all systems are linked to a central administrative centre to and from which strategic and marketing information can be transferred. The approach is particularly attractive in cases where distribution (of goods) is distributed, yet allocation of stocks against urgent orders must be carried out centrally, within one computing system which is responsible for this function.

4.2 'User sovereignty'

It is claimed that many users now using distributed systems have accepted computerisation more readily than they would have accepted a remote centralised on-line system totally under the control of a d.p. department. This results from the provision of local processing facilities partly under the control of local line management; which considerably increases the confidence of line staff and management.

4.3 Improvement in efficiency

The small computers found within the distributed systems are normally dedicated, at least for a period of time to one task only. They thus have a much higher efficiency in terms of cost per unit of work done than large multi-programming multi-tasking centralised systems. The situation is shown schematically in Figure 4, which shows that a dedicated small computer spends a far higher proportion of the available processor time processing user tasks than performing housekeeping operations.

4.4 Flexibility

As highlighted later in this chapter, many organisations find that their information processing requirements change with time. A distributed approach allows rapid reconfiguration, since the modular satellite systems can be easily moved to another location, and reprogrammed when required. Such subsystems can be replaced, enhanced, redeployed or even totally removed without substantially disturbing the remainder of the system. The risk of obsolescence

Typical activity time plots for processors within alternative configurations

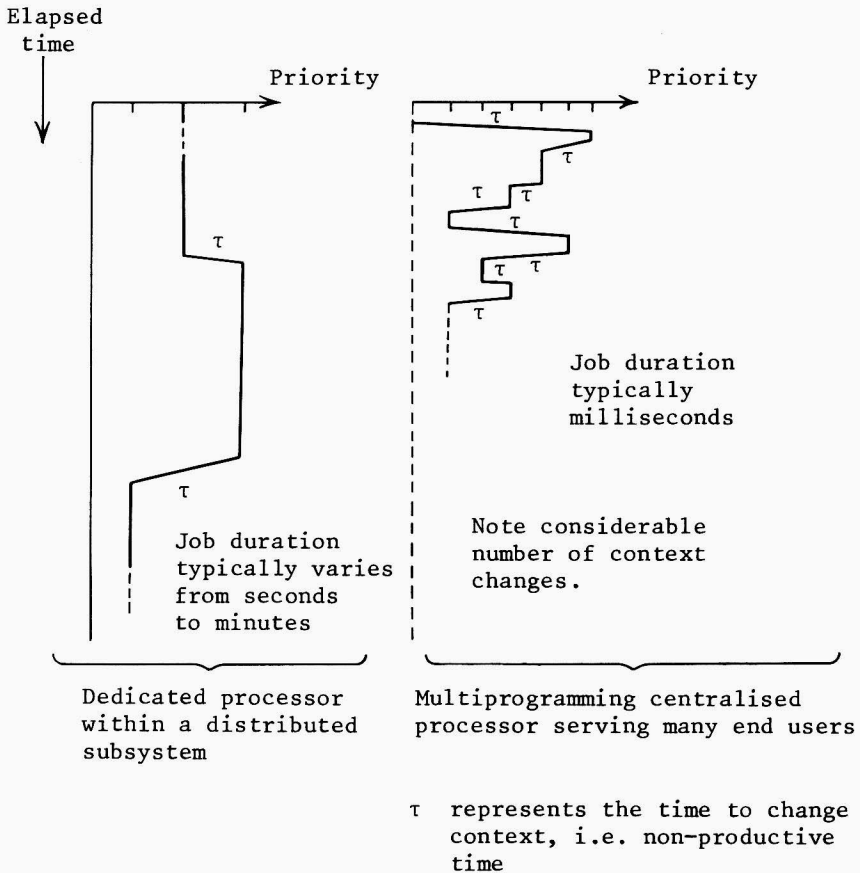


Fig.4. The processors within distributed systems change context less often than those of centralised systems, leading to increased processor efficiency and hence effective use of resources.

associated with many monolithic, expensive, totally centralised systems is therefore greatly reduced, and in some cases removed entirely.

4.5 Resilience and reliability

Distributed systems can be designed to have a higher availability at lower cost than centralised systems. Resilience can be provided either by means of off-line standby procedures which involve