NCC
The National Computing Centre

REVIEWING YOUR DATA TRANSMISSION NETWORK

This book is concerned with conventional data networks ('wide area' networks, as distinct form 'local area' networks). The review is seen as an exercise to help you to 'know your network' — in terms of how users view the system, how the network is performing internally, and how much it is costing. It is emphasised that a review helps to identify areas of inefficiency and excessive cost, and provides a firm foundation for further development of the network.

Attention is given to the parameters which indicate the user's perception of the quality of the service, to the need for an analysis of the traffic flowing in the network, and to an examination of the associated costs. Appendices are included to provide information on the services and equipment that are available.

Peter Scott, Manager of the NCC Communications Division, is author of *Modems in Data Communications* and co-author of the *Guide to British Telecom Services*.

Published by NCC Publications
ISBN 0 85012 326 7

NCC

The National Computing Centre

Reviewing Your Data Transmission Network

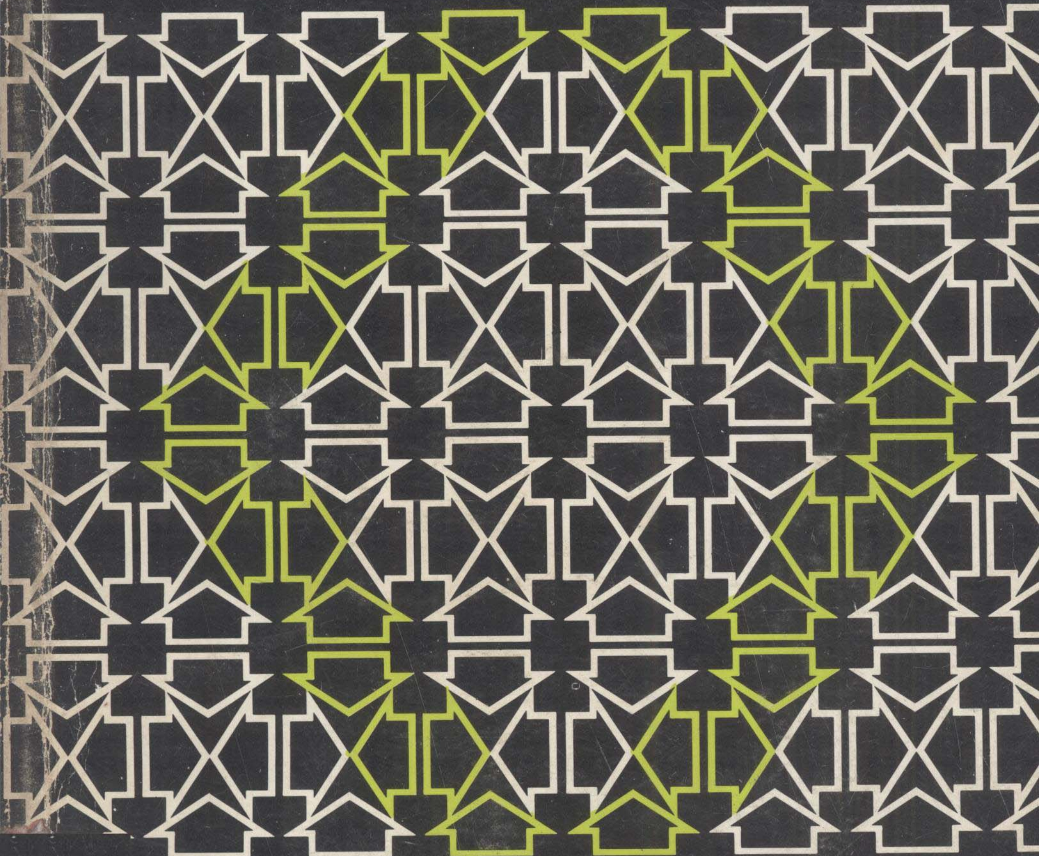

P R D Scott

Acknowledgements

Much of the material upon which this book is based, together with the appendices, was provided by PACTEL Ltd.

Thanks are also due to individuals from the following organisations who provided guidance at the early planning stages of the project and commented on a preliminary draft of the book:

Christian Salveson Ltd
ICI Pharmaceuticals
Northumbrian Water Authority
Reed International Ltd

The Centre acknowledges with thanks the support of the Computers, Systems and Electronics Requirements Board for the project from which this publication derives.

Reviewing Your Data Transmission Network

P R D Scott

PUBLISHED BY NCC PUBLICATIONS

British Library Cataloguing in Publication Data

Scott, P. R. D.
Reviewing your data transmission network.
1. Data transmission systems
2. Computer networks
I. Title
001.64'404 TK5105.5

ISBN 0-85012-326-7

First published in 1983 by:

NCC Publications, The National Computing Centre Limited, Oxford Road, Manchester M1 7ED, England.

Typeset in 11pt Times Roman and printed in England by UPS Blackburn Limited, 76-80 Northgate, Blackburn, Lancashire.

ISBN 0-85012-326-7

Contents

	Page
Acknowledgements	
1 Introduction	9
General	9
Definitions	12
2 Quality of Service	13
Introduction	13
System Availability	15
Measuring Availability	16
Improving Availability	18
Response Time	21
Terminal Access Time	21
Input Time	21
Output Time	23
Terminal Response Time	23
Measurement of Terminal Response Time	26
Setting Targets for Response Time	27
Improving Response Time	29
Presentation of Response Time and Availability Measurements	33
Summary	33
3 Traffic Analysis	35
The Need for Traffic Analysis	35
Hardware-Independent Traffic Analysis	35

Hardware-Dependent Traffic Analysis 37
Methods of Traffic Measurement 39
Software Packages 42
Improving Network Utilisation 43
Network Growth 44
Summary 46

4 Costs 47

Introduction 47
Network Costs 47
Ancillary Costs 51
Costing Techniques 51
Opportunities for Cost Reduction 51
Staff Costs 59
Summary 59

5 Conclusion 61

General 61
Security 63
Long-Term Plans 67
Summary 67

Appendix

A Public Switched Networks 69
B Network Control 73
C Network Hardware 85
D Bibliography 107

Index 109

1 Introduction

GENERAL

The 1970s were the decade of the computer network. At the beginning of the decade on-line networks were only for the adventurous, but by the end of the decade they were commonplace, and the adventurous had moved on to discover the uncharted territory of 'local area networks'.

This book is about conventional data networks or 'wide area' networks as they are sometimes called in order to differentiate them from their newer counterparts. Essentially the review described is an exercise in how to 'know your network', in terms of how the users view the system, how the network is performing internally, and how much it is costing.

Carrying out a review requires dedicated effort and consumes time and resources. Many organisations will first have to acquire the measurement tools which allow such a review to be undertaken, and so results cannot be made available overnight. The benefits can, however, be considerable. Not only will a review help to identify areas of inefficiency and excessive cost, but it can also provide a firm factual foundation for further development of the network.

It is possible to identify several distinct phases in the development of a network:

— assessment of need;

— design;

— implementation;
— post-implementation review;
— operation;
— enhancement;
— replacement.

It takes one or two years to build a network following the initial design, and most UK data networks are now in the operation/enhancement phase. The introduction of digital leased circuits could, however, have a major impact on today's data networks; the relatively low cost of high-speed digital transmission may encourage redesign and premature replacement of existing networks. Nevertheless, replacement in most cases will be two or three years away, and much can be done in the meantime to improve the efficiency of the present network.

This book describes a review to be carried out during the operational phase of a network's life. It is not a post-implementation review, nor is it a complete network redesign. It is rather an assessment of how the network is currently meeting the demands placed upon it.

In many cases, current demands differ considerably from those envisaged at the design stage. Few networks grow in a completely ordered manner, in strict conformity to a predefined strategy. Haphazard growth in a variety of directions is usually superimposed upon any planned expansion pattern, and this results in sub-optimum solutions. Also, rapid technological progress means that yesterday's state-of-the-art communications hardware may no longer be cost-effective in today's network.

These two facets of changing demand and technological progress provide opportunities for making savings. The aim of the review is to provide a framework through which the network manager can identify these opportunities and take appropriate action.

The post of network manager as such does not exist in many organisations. Typically, a company with a data communications network of any size will have a network controller, responsible for day-to-day network operation, who reports to the operations manager. The one lacks the remit and the other the time to manage the network in the true sense of the word; ie to:

— oversee day-to-day operation;

— plan the future course of the network in terms of enhancement, expansion, modification and replacement;

— evaluate the network's effectiveness;

— control expenditure.

Yet with the growing importance of communications, it is vital that an organisation's communications systems are properly managed. In the absence of forward planning, expedients which solve a short-term problem can become impediments to longer-term development.

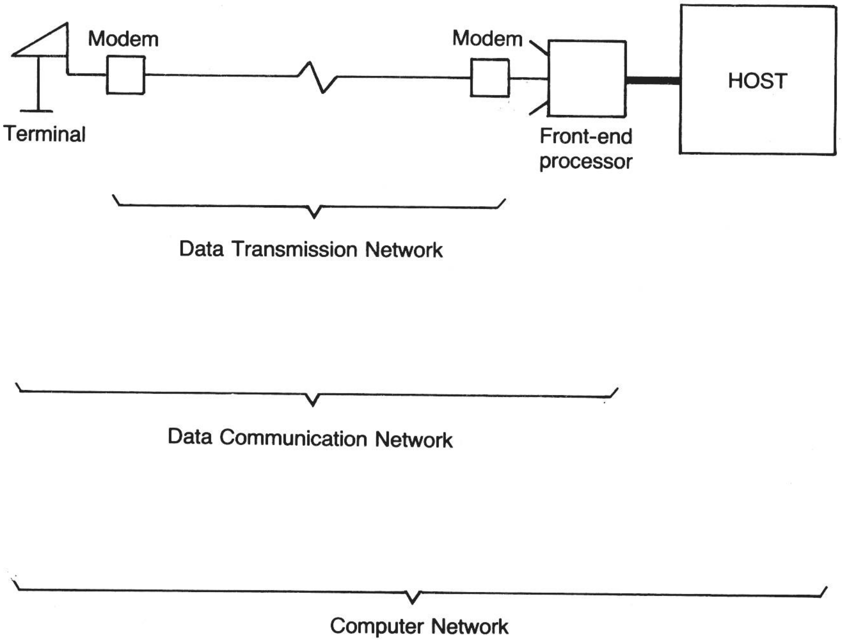

Figure 1.1 Terms and their Usage

It is unfortunate that certain equipment suppliers have coined the term 'network management' to describe their fault diagnosis and control systems. These 'network management' systems are essentially tools to use in solving daily operational problems. True network management, much more a strategic activity looking towards long-term goals and seeking the best route forward, demands both managerial and technical skills.

In summary, reading this book will not on its own make for a better network. A network review requires time to be set aside and probably some expenditure on equipment if real benefit is to be achieved.

DEFINITIONS

Definitions are given, because of the different meanings attached to words like 'network'. Figure 1.1 shows how the following terms are used:

— data transmission network;

— data communications network;

— computer network.

The *data transmission network* refers to lines, modems, concentrators and the like, and is bounded by the digital (V.24) interfaces of the modems at either end of a circuit.

The *data communications network* encompasses all the other hardware and software dedicated to the communications function, including the terminals, front-end processor and any teleprocessing software in the mainframe.

The *computer network* is everything, including the host processor.

Where 'network' on its own is used, the context will normally make clear which type of network is being discussed. This book is basically concerned with the data transmission network, although inevitably it overlaps into other areas on occasions.

2 Quality of Service

INTRODUCTION

This chapter deals with the parameters which indicate the *user's* perception of the quality of service offered by an on-line system. (It does not cover technical performance parameters such as line utilisation which, although of interest to the dp department, have little relevance to the end user. They are dealt with in Chapter 3.)

It is commonly observed that the user's view of a service differs from the provider's view of that same service, and this is one of the main causes of disharmony between user and provider. One reason for this difference in perceived quality is that user and provider employ different criteria against which to judge the service. Measures employed by the service provider are often meaningless to the user, while user measures may appear imprecise and vague to the service provider.

The dp department operating an on-line system finds itself in a dual role as both a *user* of British Telecom's communications services, and a *provider* of service to its own network users. This ought to qualify it to be the ideal service provider but, as many end users can testify, this is not always the case.

Disharmony is best avoided by the dp department, in the context of available resources, agreeing with each user department:

— what parameters should be measured in order to give an objective indication of the quality of service;

— what target figures should be set for each parameter.

A fundamental assumption is, of course, that the type of service provided meets the user's requirements.

In an on-line system, the two parameters of terminal response time and availability are the most commonly used indicators of quality of service. They are meaningful to the end user, fairly easily measured by the dp department, and give a good overall picture of system performance. They are not the only parameters which can be measured, but because of their universality they merit detailed discussion. It must be remembered, however, that there are many other aspects of quality of service which do not lend themselves to simple measurement, such as:

— the 'friendliness' of the dialogue;

— the 'friendliness' of the user-support facilities;

— the 'friendliness' of the switch-on/log-on procedure;

— documentation;

— confidentiality;

— fault reporting and clearing procedures.

All these contribute to a user's view of the system.

An organisation needs channels of communication via which user's views on these other aspects of service can be passed back to the dp department. Possible channels include:

— regular dp/user department review meetings;

— user satisfaction surveys;

— visits by dp staff to user sites;

— visits by users to computer centre;

— on-line reporting system;

— 'suggestions box' approach, possibly on-line.

The remainder of this chapter is devoted to a discussion of availability (and the related parameter reliability) and terminal response time. It covers:

— how they are defined;

— how they are measured;

— how they can be improved.

SYSTEM AVAILABILITY

Availability is the proportion of time (typically the proportion of the working day) that a system is capable of performing its designed job for the end user. Availability can be expressed as:

$$\frac{\text{MTBF}}{\text{MTTR} + \text{MTBF}}$$

where MTBF is the Mean Time Between Failure (and is a measure of the reliability of the system) and MTTR is the Mean Time To Repair the failure. MTBF and MTTR are usually measured in hours; only operational time (ie hours when the equipment is supposed to be available) are counted.

The above expression shows up the relationship between reliability and availability. A system can be very reliable (high MTBF) but if on the rare occasion when it does fail it takes a long time to repair, the overall availability will be low. Conversely, an unreliable system can give good availability if the repair time can be made very short.

Consequently it is important for the dp and user departments to agree upon targets for both reliability and availability. This involves deciding:

— the hours of day during which the service should be operational;

— planned down-time for maintenance, where applicable;

— how long the system can be unavailable before business is seriously affected;

— an estimate of likely fault frequency, based on past experience;

— an allowance for the unforeseen.

In some organisations the on-line system is crucial to the minute-by-minute running of the organisation's business, whereas, in others, short periods of down-time can be tolerated. Consequently it is well nigh impossible to lay down norms at which an organisation should aim. Some networks come very close to 100% availability (eg greater than 99.5%), but at this level the cost of achieving even the slightest increase in availability can be considerable.

It is also important to define exactly what is being made available. Is it the availability of:

— the required application programs at all user terminals?

— the required application programs at some proportion of the user terminal population?

— the required application programs at specified user terminals?

— the required application programs on the host machine, with no reference to user terminals?

Ideally the availability of the required applications at each and every user terminal should be known, although practically something less comprehensive than this may have to be accepted. Table 2.1 shows possible reasons for *un*availability, and actions that can be taken.

MEASURING AVAILABILITY

In a multi-component system like a network, it is often easier to calculate availability figures than to measure them directly. Availability of service to the end user (As) depends on the availability of the host (Ah) and the availability of a communications path from the user's terminal to the host (Ac). Hence:

$$As = Ah \times Ac$$

The ***host system availability*** is readily found from the daily log. The ***communications path availability*** can be found easily if the network utilises intelligent data communications equipment in conjunction with a network management and control system (see Appendix B). Failing this, it will be necessary to work out the availability of the whole communications path from records of the