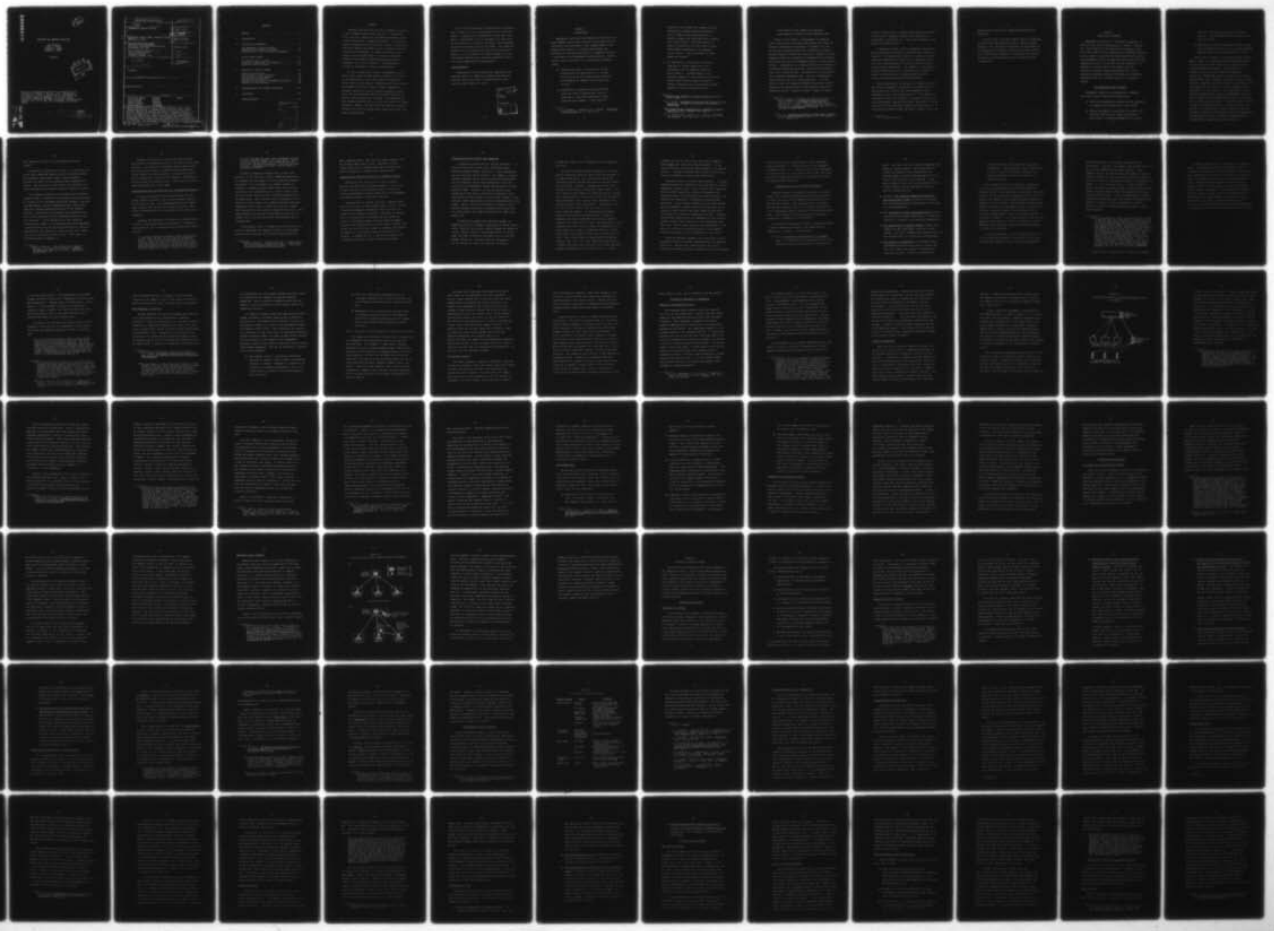
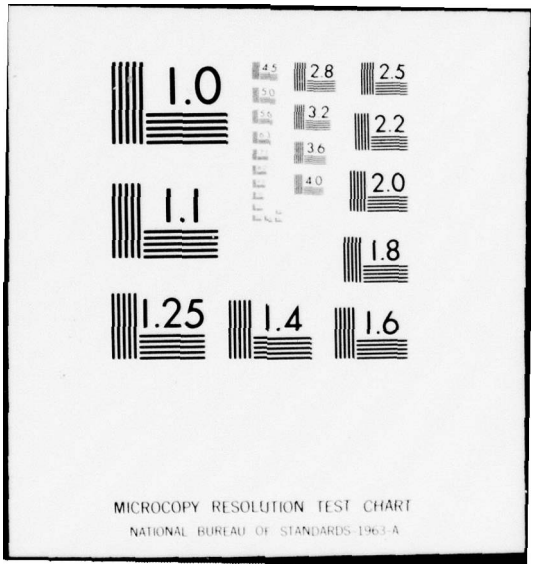
AD-A038 082 WHARTON SCHOOL OF FINANCE AND COMMERCE PHILADELPHIA P--ETC F/G 9/2
CHARGING FOR COMPUTER SERVICES.(U)
FEB 77 D BERNARD, J C EMERY, R L NOLAN N00014-75-C-0462
UNCLASSIFIED 77-02-06 NL

1 OF 2
ADA038082




1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD A038082

(12)

CHARGING FOR COMPUTER SERVICES

Dan Bernard
James C. Emery
Richard L. Nolan
Robert H. Scott

77-02-06

DDC
RECEIVED
MAR 22 1977
C

This is a preliminary version of a monograph to be published by EDUCOM in Spring, 1977. Reproduction in whole or in part is permitted for any purpose of the United States Government. All other rights reserved by EDUCOM, The Inter-University Communications Council, Inc., P. O. Box 364, Princeton, New Jersey 08540.

DDC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 77-02-06
2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle): Charging for Computer Services
5. TYPE OF REPORT & PERIOD COVERED: Final rept.
6. PERFORMING ORG. REPORT NUMBER: 77-02-06
7. AUTHOR(S): Dan Bernard, James C. Emery, Richard L. Nolan and Robert H. Scott
8. CONTRACT OR GRANT NUMBER(S): N00014-75-C-0462
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Decision Sciences Department, University of Pennsylvania/Wharton School, Philadelphia, PA 19104
10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: Technical report
11. CONTROLLING OFFICE NAME AND ADDRESS: Office of Naval Research, Information Systems, Arlington, Virginia 22217
12. REPORT DATE: 2/77
13. NUMBER OF PAGES: 119
14. MONITORING AGENCY NAME & ADDRESS (If different from Controlling Office): 118 p.
15. SECURITY CLASS. (of this report): Unclassified
15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report): Unlimited
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)
18. SUPPLEMENTARY NOTES
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Information Systems	Charging	Control
Data Processing	Pricing	
Computer Centers	Budgeting	
Computer Services	Economics	
Cost Allocation	Accounting	

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Charging internally for the use of computer facilities is now a common organizational practice. A charge-out policy can play a major role in promoting effective and efficient utilization of computing resources. In practice, charging all too often fails to have a significant beneficial impact, and indeed can be a source of tensions and user dissatisfaction. A charge-out system is most likely to be successful when it is based on an understanding of the purposes underlying charging and the requirements for it to be effective. This monograph is intended to contribute to such an understanding.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE S/N 0102-014-6601

408 957

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

CONTENTS

PREFACE ii

1. INTRODUCTION 1

2. OBJECTIVES IN CHARGING 6

The Management Control Problem 6
Functions of a Charge-Out System 8
Limitations of a Cost Allocation Approach 17

3. POLICY-LEVEL ISSUES 22

To Charge or Not to Charge 22
Alternative Approaches to Budgeting 30
Financial Control Options 48

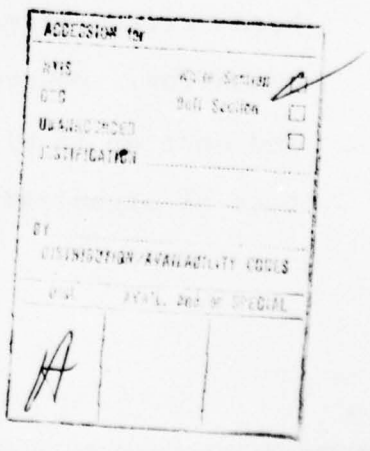
4. DESIGN OF A PRICING SCHEME 58

Underlying Principles 58
Establishing a Price Structure 70
Alternative Pricing Methods 78
Pricing Service Priority 86
Charging for Systems & Programming Resources . . 90
Special Pricing Problems 95

5. ADMINISTRATIVE AND CONTROL PROVISIONS 102

6. CONCLUSION 107

FURTHER READING 111



PREFACE

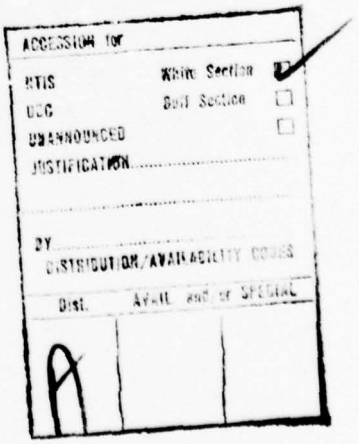
Charging internally for the use of computer facilities is now a common organizational practice. A charge-out policy can play a major role in promoting effective and efficient utilization of computing resources. In practice, however, charging all too often fails to have a significant beneficial impact, and indeed can be a source of tensions and user dissatisfaction. A charge-out system is most likely to be successful when it is based on an understanding of the purposes underlying charging and the requirements for it to be effective. This monograph is intended to contribute to such an understanding.

We view a charge-out system as a management control device which, like all such devices, needs to be tailored to the objectives it is to serve and the circumstances within which it will operate. Therefore our discussion does not attempt to specify a single 'correct' approach to charging; instead, we lay out design principles, alternatives, and considerations which will provide guidance in dealing with any particular situation. The monograph has been written primarily for use within institutions of higher education. However, the discussion has been kept general, and most of it will be found equally applicable within any type of organization.

One of the difficulties in dealing with the charging area is that it is one in which there are few clear-cut answers. Charging can be approached in many ways, and there is considerable lack of agreement as to the nature of the problem, let alone the solutions! This monograph therefore inevitably reflects the authors' own viewpoints. Readers may well disagree with some of our conclusions. However, if in the process we provoke thought on issues which might otherwise have remained unexamined, then we will have accomplished our main purpose.

Acknowledgement

Preparation of this document was supported by the Planning Council on Computing in Higher Education and Research, and by the Office of Naval Research under contract number N00014-75-C-0462.



Chapter 1

INTRODUCTION

Management control has been defined as "the process by which managers assure that resources are obtained and used effectively and efficiently in the accomplishment of the organization's goals."[1] Assuring effectiveness and efficiency in the provision and use of computing resources is one aspect of the management control problem. Moreover, several factors have made it an increasingly important aspect:

- Computing now plays a critical role in most organizations of significant size. In many cases most of the key organizational functions are dependent on computer-based systems and services.

- Expenditure on computing has grown with the range and scale of applications, and now represents a significant proportion of most organizational budgets. A 1968 survey of

---

[1] R. N. Anthony, J. Dearden & R. F. Vancil: Management Control Systems, R. D. Irwin, 1972

industrial firms found that computer-related expenditures averaged around one percent of sales;[1] a more recent sample provided a similar average, with individual cases ranging up to four percent.[2] Data published for a group of universities in 1971 showed an even higher level of expenditure on computing in this sector -- over two percent of total budget on average.[3]

- The importance of effectively controlling computing is further emphasized by the 'Parkinson's Law of computing' noted by Martin Robbins -- workload tends to fill available capacity.[4] Experience has shown that without some effective means of control, computer resources have a particularly strong tendency to be used ineffectively and

---

[1] Unlocking the Computer's Profit Potential, McKinsey & Co., 1968

[2] R. L. Nolan: Management Accounting and Control of Data Processing, National Association of Accountants, forthcoming

[3] The Financing and Organization of Computing in Higher Education: 1971, EDUCOM, 1971, page 8

[4] M. D. Robbins, W. S. Dorn & J. E. Skelton: Who Runs the Computer? Westview Press, 1975, p. 63

inefficiently, while demand for computing seems capable of growing without apparent limit.

Studies of the pattern of development of computing within organizations have shown that charge-out systems are typically introduced by management in response to a perceived need to improve control over computing activities.[1] The main objective behind this decision is normally to bring computing within the framework of responsibility accounting, which forms the basis for management control in most organizations. Charging has by no means solved all the problems of controlling computing activities, and indeed has sometimes given rise to some new ones. Nevertheless, it has proved to be an important management tool. The use of charging has grown steadily, and is now normal practice: in a 1975 survey by the Cost Accounting Standards Board over seventy percent of the organizations polled operated some form of charge-out scheme.[2] Moreover, Nolan's 'Stages of Growth'

---

[1] See R. L. Nolan: "Managing the Computer Resource: A Stage Hypothesis," Communications of the ACM, Vol. 16, No. 7, (July 1973), pp 399-405, and C. F. Gibson & R. L. Nolan: "Managing the Four Stages of EDP Growth," Harvard Business Review, Vol. 52, No. 1, (January-February 1974)

[2] D. H. Li: Accounting for Costs of EDP Service Centers, Cost Accounting Standards Board, 441 G St., Wash., D. C. 20548, 1975

studies[1] indicate that, as organizations become more mature in their use of computers, charge-out systems play an increasingly significant role. Thus their importance seems likely to continue to grow.

A charge-out system, like any management control tool, must be designed in relation to the particular situation involved. The features best suited to a particular organization will depend on the nature of its computing activities, the sophistication of its users, and other factors unique to that organization. For example, an installation serving research users on a 'job-shop' basis probably requires a different basis of charging than one which carries out routine administrative processing.

More significantly, the design of a charge-out system must reflect management's objectives in controlling computing activities and the role that management wishes charging to play in the control process. For example, the importance of the role assigned to a charge-out system can vary considerably, ranging from simple 'funny-money' schemes that act as little more than resource rationing devices, to financially-oriented systems that play a central role in the

---

[1] R. L. Nolan, op.cit.

organization's decisions on computing expenditure and allocation.

In view of this, the next chapter begins by examining the objectives that potentially may be pursued through charging. This discussion establishes the background for the remainder of the paper, in which our primary concern is to show how the features of a charge-out system can be tailored to suit any particular set of objectives and circumstances.

Chapter 2

OBJECTIVES IN CHARGING

Management's objectives in charging for computer services may vary from one organization to another. Nevertheless, these objectives are typically all related in some way to the control of the organization's computing activities. In this chapter we consider the broad objectives associated with management control of computing, and the functions that charging can fulfill in relation to these objectives. In addition, we discuss some of the dangers that are inherent in viewing viewing charging purely as a cost allocation mechanism, without considering the underlying objectives that this allocation process is intended to achieve.

The Management Control Problem

Management control of an organization's computing activities centers around three related tasks:

- Achieving an appropriate balance between expenditure on computing and expenditure in other areas.
- Allocating computing expenditure to the most cost-effective combination of computing resources -- particularly in balancing expenditure on central

facilities, distributed systems, and external services -- and ensuring that these resources are used efficiently.

- Deploying the available computing resources among users and application areas in a way which provides maximum benefit to the organization as a whole.

These tasks are common to any organizational resource. However, there is an additional associated problem which arises specifically from the particular characteristics of shared computing resources. Such resources are normally set up as a central function supporting several user departments within an organization or major organizational unit. However, unlike most central support functions, computing typically has major significance within the organization in terms of its cost, complexity, and importance as a key resource. This significance tends to generate strong pressures from user departments to gain direct control of their computing activities, on the same basis as other major resources. The same factors also tend to make such decentralization of control desirable. The problem is to decentralize control effectively when the resources involved are shared between the user departments, and when many key issues relating to computing must be decided from an organization-wide perspective.

Functions of a Charge-Out System

Consider now the functions of charging in relation to these tasks. All charge-out systems consist of two interdependent components:

- a budgeting process through which the organization plans the provision of computing resources and determines their internal allocation.
- a pricing scheme that measures, and provides a basis for controlling, users' consumption of these resources.

The combination can be used to support all the objectives discussed above. Specifically, it can:

- provide a means of allocating computing resources among competing users.
- promote effective and efficient use of these resources.[1]

---

[1] Our distinction between effectiveness and efficiency is essentially that between 'doing the right things' and 'doing things right.' In this context, effectiveness is a function of the choice of applications for computer resources, while efficiency is concerned with the cost of the resources used in these applications. Clearly both are involved in obtaining maximum benefit from computing expenditure.

- promote cost-effective provision of computing services.

- permit decentralization of decisions associated with computing activities, particularly those relating to the scale and allocation of computing expenditures.

These functions are discussed further below. We would point out, however, that in each case we are discussing functions that charging can perform. The actual functions performed by any particular charge-out system will depend on management's view as to the role charging should play in the overall management control process, and on the extent to which the system has been designed to carry out effectively its intended functions.

Allocation of Computing Resources

Unless an organization is in its early stages of computer use, the potential demand for computing resources typically far exceeds the capacity which is available or which could be justified in cost-benefit terms. Thus there is a need for some method of allocating these scarce resources among competing users. As Nielsen has pointed out, "if resource allocation is not done explicitly, it will be

done implicitly; there is no such thing as "no allocation."[1]

Allocation is needed at two levels -- to decide how use of the total computing capacity will be distributed among the alternative application areas, and to determine priority for access to the resources under congested conditions. The objective at both levels will generally be to obtain maximum overall benefit from the resources available.

By providing a mechanism for planning and controlling computer usage, a charge-out system enables the distribution of resources to be determined explicitly through a conscious (and presumably rational) budgeting process, rather than implicitly through competing pressures on the computer facility. This process also provides each budgetary unit with a clear indication of the computing resources it will have available. In the simplest case, the charge-out system can act simply as a rationing device which implements a central decision on the desired distribution of resources. Alternatively, the charge-out system can also play a major role in the allocation decision itself. (This choice is discussed further in Chapter 3).

---

[1] Norman R. Nielsen: "The Allocation of Computer Resources--Is Pricing the Answer?" Communications of the ACM, Vol. 13, No. 8 (Aug. 1970), pp. 467-474.

Charging can also play a role in handling the lower-level problem of determining priority for access to the shared resources. By setting up a charging structure where prices vary with service level, the organization can effectively set up a market for access priority, in which users can be left to decide their service level themselves, depending on the price they are prepared to pay. (This point will also be discussed further at a later stage).

Promoting Effective and Efficient Use of Computing Resources

A charge-out system limits the availability of computing, thereby turning it into a scarce resource from the user's point of view. The charging mechanism also communicates to users the cost of the resources involved. These effects encourage users to be cost-effective in their use of computing.

Charging can influence user behavior at several levels. Firstly, it discourages waste of resources. An excellent illustration of this is provided by an incident described by Nielsen:

> "...an acute shortage of magnetic tapes had developed at the Campus Facility of the Stanford Computation Center. Users were reserving tapes at an ever increasing rate; yet they were releasing very few tapes. Pleas were made for users to release all tapes that were not absolutely essential for their work, but this met with almost no response. Finally, in an attempt

> to cover the mounting tape costs, management decided to levy a nominal charge of $1 (tape) month. On the first day that the charge was instituted more than 1/3 of the 'absolutely essential' reserved tapes were released. Management has been a firm believer in pricing ever since."[1]

Secondly, charging encourages users to make cost-effective use of resources in their computing applications. For example, one user might be led to design his programs efficiently so as to minimize his charges; another to eliminate a regularly printed report whose value did not justify the cost of producing it; a third to make a cost-effective choice between on-line and batch operation for a new application. Note, however, that if charging is to be effective in this respect, the computer pricing scheme must adequately reflect costs, so that user efforts to reduce their charges do in fact result in improved efficiency. More significantly, the user must understand the basis for his charges if he is to be able to seek intelligently ways of reducing them.

At the highest level, a charge-out system can assist in effective deployment of computer resources by encouraging users to select the most cost-effective applications for

---

[1] Norman R. Nielsen: "Flexible Pricing: An Approach to the Allocation of Computer Resources," Proc. AFIPS Fall Joint Computer Conference, 1968.

their computing funds. Note that this again requires that the pricing scheme reflect costs. Moreover, it also assumes that users go through some form of rational cost-benefit evaluation in choosing their applications.

Promoting Cost-Effective Provision of Computer Services

Charging for the services of a central computer facility (or facilities) provides guidance as to the cost of these services, and can thereby assist in the cost-effective deployment of funds between the central facility, distributed resources, and outside services.

Charging also can promote efficiency, together with quality of service, at the computer center itself. Firstly, it provides management information which can assist in control of the center's operations. Secondly, by placing users in the position of customers who are paying for service, it generates market pressures that can promote efficiency and responsiveness to local needs. Thirdly, the center can use the pricing scheme to influence the pattern of demand, so as to achieve balanced utilization of resources: an example of this is the use of discounts for overnight processing to smooth installation loading.

Decentralization of Control Over Computing

A charge-out system exercises control indirectly -- it leaves the decisions to the users, influencing these decisions through economic incentives. Clearly management can also exercise control directly. For example, a central steering committee can be used to decide what applications will be developed, rather than leaving this to the discretion of user departments within the constraints of their available funds. Similarly, at the operational level, scheduling of jobs may be carried out according to fixed rules based on user category, or left to the Computer Center Director, rather than using a priority-differentiated pricing scheme. Again, efficient utilization of resources can be pursued through restrictions on program run-time and core size, disk quotas, over-night printing of all large-scale output, etc.

In practice an organization typically manages its computing activities through a combination of direct administrative controls and indirect control through the charge-out system. However, there is considerable choice as to the range and scope of decision areas that will be controlled through charge-out. The more extensive management's use of charging as a basis for control, the greater

freedom this leaves users to manage their own computing activities.

Any form of charging scheme eliminates, or at least greatly reduces, the need for restrictions and controls at the operational level. This already gives users substantially greater autonomy than would be possible under a free-use policy. Moreover, charging can be used to decentralize decisions at considerably higher levels: users can be given a total budget for computing and left to make their own decisions as to how these funds will be deployed among alternative uses, based on their judgment of the relative value of these alternatives. Furthermore, users can also be given control over the actual size of their computing budgets. By integrating the computer charge-out system with the organization's financial budgeting and control mechanisms and treating computer charges on the same basis as 'real' cash expenditures, each budgetary unit can be left to choose the proportion of its total budget to be spent on computing, based on its own perception of the tradeoff between this and other uses of funds. Thus, rather than the organization determining centrally the size of the total computing budget and its internal allocation, these decisions can be made implicitly through a combination of individual choices. Finally,

charging can also be used to give users more freedom to choose where they spend their computing funds -- at the central facility, on their own mini-computer, or on outside services -- thereby decentralizing decisions on the distribution of expenditures between these alternatives.

Decentralization can provide major benefits. It brings computing within the framework of the organization's overall management control system. Direct controls over computing, which are often inflexible, unwieldy, unresponsive to individual needs, and distorted by political pressures, are replaced with impersonal, economically-oriented market mechanisms. Given rational computer pricing, these mechanisms will generally promote effective and efficient use of resources, while leaving individual organizational units with control over their own computing activities. Moreover, transferring responsibility for computing to users tends to greatly increase their interest in, and involvement with, this area. This not only improves the effectiveness of computer use, but also increases pressures on the computer facility to provide a service responsive to user needs.

These are powerful arguments for giving users as much freedom of choice as possible regarding their own computing activities. However, the extent to which this can effectively be achieved through a charge-out system will depend

on such factors as the organization's overall management control philosophy, the sophistication of line management as computer users, and the degree to which senior management is satisfied that decisions taken by users within the framework of the charge-out system will be consistent with overall organizational objectives.

Limitations of a Cost Allocation Approach

Charging schemes are frequently regarded as straightforward cost allocation mechanisms whose basic objective is to recover the costs of the computer center from users in some equitable manner. This approach fails to recognize that the cost allocation process is normally motivated by a set of management control objectives of the type discussed above. These underlying objectives must be explicitly considered if charging is to be effective.

The problem of designing an effective charge-out system is greatly complicated by some of the complexities of computing as a resource. In particular:

- Due to the economic characteristics of computer costs -- in particular the high proportion of fixed and joint costs and the presence of economies of

scale -- the true 'cost' of a specific computer task cannot be clearly defined. Fairly arbitrary decisions are required in the construction of a pricing scheme, and these decisions need to be related to the objectives to be achieved. In the process, implicit costs such as the opportunity cost of idle capacity must be considered as well as explicit accounting costs.

- Because of the technical complexity involved in measuring resource usage, the desire for accuracy must be balanced against considerations of cost and user understandability.

- With the variety of users and applications typically associated with a large installation, a single price structure is unlikely to be universally appropriate.

- The significance of computer charges in many users' budgets can make them highly sensitive to these charges. This gives computer cost allocation many of the characteristics of a transfer pricing problem.

- The presence of competition in the form of small departmental facilities and outside services may have to be taken into account in designing the charging structure, in order to encourage the

desired distribution of expenditure among these alternatives. The problem may also be complicated by sale of services to outside customers, for which the pricing objectives are likely to be rather different.

These complexities have two important consequences. First, the straightforward cost allocation approach to charging provides little guidance in designing a charge-out system to act as an effective management control tool. For example, the sophisticated pricing techniques employed by commercial service bureaux, such as output-related pricing, quantity discounts, and surcharges for use of specialized software, can greatly improve the effectiveness of an internal charging scheme. (The use of these techniques is discussed in Chapter 4). Yet their potential value is only apparent if one thinks in terms of charging as a control mechanism; they seem quite irrelevant from the viewpoint of simple cost allocation.

A second consequence is that principles inherent in the cost allocation approach, such as full cost recovery and equitable charging, are not necessarily appropriate

where the objective is to promote efficiency and effectiveness.[1] Consider, for example, cost recovery policy in the case when demand does not match capacity. Since computer capacity can generally only be changed in large increments, periods of over or underutilization are bound to occur as a result of configuration changes or fluctuations in demand. If a cost recovery policy is rigidly adhered to, prices are driven up during periods of underutilization (thus further reducing demand), while the opposite effect occurs during overutilization.[2] Given that management is normally concerned with matching usage with capacity and maintaining stability of prices, it may be preferable to accept a deficit or surplus of charging revenue in order to achieve these objectives.

---

[1] Legal requirements may limit an organization's freedom to diverge from these principles. A particular case is that of organizations with installations which derive a significant proportion of their income from government contract work. Federal policy requires that charges for such work be based on rates designed to recover the full cost of the installation over each accounting period, and applied on an equal basis to both government-funded and other users. This policy acts as a severe constraint on the organization's ability to manipulate prices in the interests of effectiveness and efficiency. (For a further discussion of this problem, see H. Kanter, A. Moore & N. Singer: "The Allocation of Computer time by University Computer Centers," Journal of Business, Vol. 41, No. 3 (July 1968, pp 373-384).

[2] Kanter, Moore & Singer (op cit) discuss this problem.

A further example occurs in the design of a charging scheme. The important question in designing a price structure is not whether the charges produced will be equitable (though obvious inequities are likely to be unacceptable), but whether they will have the desired influence on users. Computer charging schemes are often made very complex in order to achieve great accuracy in measuring users' resource consumption. Such complexity is expensive and often leads to charges which are incomprehensible to the user; the accuracy obtained is of value only to the extent that it is likely to materially affect users' decisions, and to the extent that the resulting improvement in effectiveness and efficiency justifies the cost involved.

Chapter 3

POLICY-LEVEL ISSUES

This chapter is concerned with issues that arise in determining the basic structure of a charge-out system. First we discuss the question of when charging is worthwhile. We then examine alternative approaches in two key areas: budgeting for computing, and financial control of the computer facility.

To Charge or Not to Charge

Cost of Charge-out Systems

Any charge-out system is expensive to operate. Billing is normally carried out through computer routines that keep track of each task's usage of system resources, generate charges for each job and report these to the user, and produce monthly accounting summaries; these routines themselves consume system resources.[1] In addition, there is a cost, both to the computer center and to users, in the

[1] An indication of the level of overhead involved is provided by Nielsen, who describes a charging system in use at Stanford University as requiring less than 1% of the system's CPU cycles (See N. R. Nielsen: "Flexible Pricing - An Approach to the Allocation of Computer Resources," Proc. AFIPS Fall Joint Computer Conf., 1968, p 531).

administrative effort required for budgeting, opening and closing charge accounts, dealing with billing errors, etc.

Apart from these obvious costs it has to be accepted that, as with all control mechanisms, charging is liable to produce dysfunctional side-effects. In other words, due to the imperfections inevitable in any practicable system of control, there are bound to be cases where charging induces counterproductive behaviour.[1] These dysfunctions can be minimized through careful design and management of the charge-out system, but they will never be entirely eliminated.

One particular dysfunction which is often advanced as an argument against charging is its tendency to inhibit full utilization of capacity. Whenever the installation is left with spare capacity because potential users do not have the requisite budget available or because they are not prepared to pay the normal processing charges, there is an opportunity

---

[1] For a general discussion of dysfunctional effects in control systems, see E. E. Lawler III and J. G. Rhode: Information and Control in Organizations, Goodyear Publishing Co., 1976.

cost to the organization.[1] Such problems may be avoidable through sufficient flexibility in budgeting and pricing, and, where appropriate, through sales to external users. Nevertheless it is not uncommon in practice to come across situations where a machine sits idle while there are users able and willing to utilize the spare resources, but inhibited from doing so by the charging policy.

The point to be made about this and the other costs of charging is that they will be accompanied by the benefits associated with improved managerial control. As Nielsen says:

> "It is true that pricing does result in more administrative work, but pricing also provides for better utilization of the available resources. While the computer may not perform any additional processing for users, the processing it does perform is more valuable to the organization. Thus the system is doing more valuable computing even if it is not doing more computing. So although pricing does involve additional costs, it returns additional benefits."[2]

---

[1] This cost may not be as great as it initially appears. A reduced workload improves the speed of service provided by the installation, and the value of this improvement to users may partially, if not fully, offset the cost of the unused capacity. In fact, where the workload is over seventy to eighty percent of capacity, the value of any increased usage is liable to be more than offset by the cost of the increased congestion produced.

[2] Norman R. Nielsen: "The Allocation of Computing Resources - Is Pricing the Answer?" Communications of the ACM, Vol. 13, No. 8 (August 1970), p. 473.

Thus the appropriateness of charging in any particular context must be judged in terms of the extent to which the benefits obtained are likely to justify the costs involved.

When Charging is Justified

Broadly speaking, the need for charging grows with the scale and complexity of an organization's computing activities.[1] When an organization is in the early stages of computer use, computing expenditures are likely to be relatively insignificant, and the potential payoff from charging correspondingly small. Management will generally be concerned more with promoting than regulating computer usage, and charging may be felt to discourage experimentation with new applications.[2] Moreover, a charge-out system relies for

---

[1] Further discussion on this topic will be found in R. L. Nolan: Management Accounting and Control of Data Processing, National Association of Accountants, (forthcoming).

[2] There are dangers in this short-term viewpoint. With computer services freely available, users are liable to develop wasteful habits and become committed to applications that may not be cost-effective. This may lead to problems with users when it becomes necessary to assert control through charging at a later stage.

its effectiveness on users making informed decisions within the framework of the computer pricing and budgeting mechanisms. With a relatively unsophisticated user community, it may be desirable to retain central control over computing activities.

As computer use grows within the organization, controls become necessary to allocate capacity and to ensure its efficient use. The normal initial approach is direct administrative control by a central steering committee and/or the computer center itself. These mechanisms may continue to be adequate as long as the level of computing expenditures remains relatively small, the range of applications limited, and the pressures on the computer center manageable. Generally, however, there comes a point where direct controls alone become inadequate. This may be indicated by symptoms such as the following:

- The computer center is continuously overloaded. Additions to capacity simply produce corresponding increases in demand. Management is alarmed at the escalation in computing expenditures, and is concerned that the resources are not being employed cost-effectively.

- Users are finding that computing plays an increasingly important role in their activities, and are frustrated with their lack of control over this area.

- Management see an increasing need to ensure that line managers become involved in the computing activities associated with their areas of responsibility, and to make them accountable for these activities.

This is the point at which charging is likely to be justified.

The charge-out system may initially play a limited role, with management retaining direct control over many key decision areas. For example, a 'funny money' charging scheme may be used to ration computer usage, with the basic decisions on the allocation of computing funds still being taken centrally. Or users may be charged for processing operations but not for system development work, the latter being controlled centrally through a computer policy committee. Experience shows, however, that as the scale and complexity of computer usage continues to grow, management has to rely more and more on the charge-out mechanism as a means of control.

Even when fully developed, the charge-out system is still likely to be supplemented with direct controls. Firstly, direct controls may be necessary where the market mechanism does not produce the desired incentives; for example, many organizations are unwilling to give users freedom to choose between internal and external sources of computing on the basis of relative costs, on the grounds that decisions made on this basis do not take account of the interest of the organization as a whole in maintaining a large, fully-utilized central computer facility. In addition, there may be areas where direct control is more cost-effective than control through the pricing and budgeting mechanisms; for example, where it is desired to reduce the number of long program runs submitted during peak periods, it may be simpler to do so by direct limits on run-time than by incorporating a premimum rate for long runs into the charging structure.

Free Service Classes

Even where charging is generally worthwhile, there may be specific categories within the overall pattern of computer use where the potential benefits from charging do not justify the costs involved. This may occur, for example, where the class of work involved has very low resource requirements, or where resource usage would not be signifi-

cantly influenced by charging. Where such categories occur and can be clearly defined, it is generally preferable to establish them as free service classes. Charges for these classes should be allocated to an appropriate organizational overhead account rather than being billed to individual users.

A particular case is that of student computing within universities. The level of resource usage for instruction-related student computing is generally quite low, while the effort involved in setting up, administering and billing student accounts for each course can be considerable. It may be preferable to eliminate direct charging of students (or at least set up student-oriented classes of use such as small compile-and-go Fortran jobs as free service classes); direct restrictions on program size, run-time, etc., can be used to restrict individual students' resource usage where necessary. One example of such a policy is that operated at Dartmouth College. Dartmouth has concluded that only a small minority of students are likely to make unjustifiably large use of computer resources under a free use policy. It has therefore adopted such a policy, exercising control over the wasteful minority through restrictions on the

total resources that can be consumed by any one student.[1]

Alternative Approaches to Budgeting

Purpose of the Budgeting Process

The computing budget normally covers the same one-year cycle as the organization's overall financial budget, though the two budgeting processes may be only partially integrated. When computing is charged out to users, computer budgeting must be carried out in two dimensions: expenditure and revenue. The expenditure budget is a line-oriented budget for the computer facility: it defines the planned total cost of the facility, and breaks these costs down into responsibility centers and categories to provide a basis for internal control of the facility. The revenue budget, on the other hand, is broken down by user area and defines each organizational unit's planned usage of computing resources (including, where appropriate, planned revenue from external users). It is the revenue budget that provides the basis for controlling computer usage through the charging mechanism.

---

[1] See A. W. Luehrmann & J. M. Nevison: "Computer Use Under a Free-access Policy", Science, Vol. 184, , 1974, pp. 957-961.

The computing budget is of key significance to all those concerned with the organization's computing activities: top management, the management of the computer facility, and computer users. For top management the budget defines the planned level of expenditures on computing resources, and also defines the planned distribution of resource usage within the organization. From the viewpoint of the computer facility management, the budget establishes the level of funding available to support the facility, and the scale of demand that will have to be met.[1] Finally, for each user area the budget provides an indication of the computing resources that will be available to it.

The end-product of any computer budgeting process must always be a pair of consistent expenditure and revenue budgets (allowing for any planned surplus or deficit of

---

[1] In general, the level of demand is expressed simply in terms of the total value of services to be provided; the computer facility then uses its own judgment to assess the likely distribution of demand between the various categories of service (batch processing, timesharing, system development, data storage, etc.), in its internal planning. However, where management needs a more reliable basis for planning and controlling the distribution of demand -- for example, between two machines operated at the same facility -- the revenue budget can be broken down to provide this.

revenue over expenditure). However, the process by which these budgets are established, and thus the role played by the various participants in making the basic allocation decisions which are embodied in the budget, can vary considerably. The key variable here is the extent to which users are left free to choose the level of their computing expenditures. We discuss below the alternatives of fully centralized and fully decentralized budgeting, and then point to schemes which compromise between these two extremes. Finally, we discuss another important variable in the budgeting process -- the extent to which users are left free to allocate their computing expenditures between central and non-central resources.

Centralized Budgeting

Under the fully centralized approach, the basic decisions on the overall level of computing expenditure and the internal allocation of computing resources are made centrally -- typically by a computer steering committee on which all interested parties are represented. The allocation is established by distributing the total planned provision of resources among users in the form of 'computer units' which can be exchanged for computer services. (The total budget allocated may be slightly higher than the revenue target to allow for some budgets not being fully

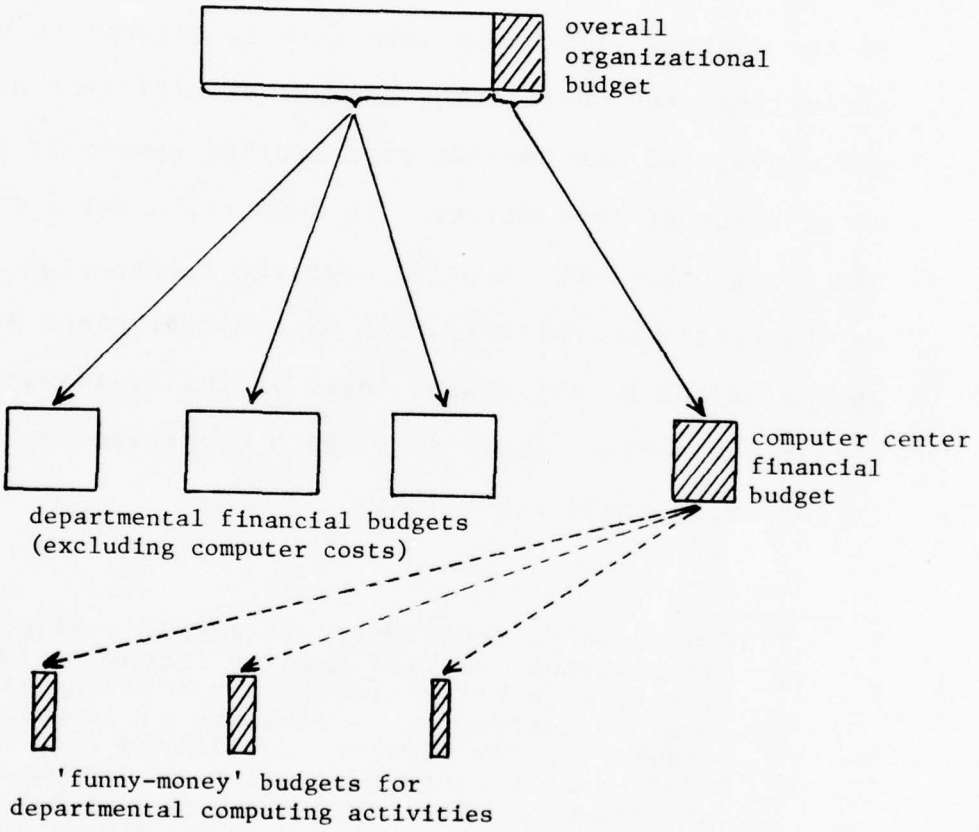
used up). Though generally expressed in dollar terms, the computer allocations can normally only be spent at the computer center and are therefore commonly referred to as 'funny-money'.

A highly simplified diagrammatic representation of the budgeting process is shown in Figure 3.1. It will be seen that with this approach, computing expenditures are treated as an overhead cost in the organization's financial budget structure; user departments' operating budgets make no direct provision for computing charges. Once the overall computing budget has been established, the process of allocating funny-money computing budgets to departments and generating charges against these budgets is entirely independent of the organization's financial system.

With a centralized budgeting process the role of the charge-out system in influencing resource allocation is strictly limited. Funny-money charging clearly performs the basic rationing function of a charge-out system and, provided the computer allocations are reasonably restrictive, will encourage users to make cost-effective use

Figure 3.1

Simplified representation of a centralized budgeting process for computing



of their resource allocations.[1] However, computing remains an essentially free resource for users, in the sense that computer use involves no sacrifice in other areas. While this encourages experimentation with new computing applications, it does nothing to restrict the level of computing resources demanded by user departments; it is left to the central allocation committee to attempt to balance effectively the conflicting demands of different user departments, and the demands of computing generally against other forms of expenditure. In summary, under a centralized budgeting procedure, control over the fundamental decisions on the extent and distribution of computer usage is retained at the policy-making level of the organization; the charge-out system simply provides a mechanism for implementing these decisions.

---

[1] The proviso is important: if a user's allocation is too generous, he will have no incentive to be selective in his choice of applications or to reduce his charges by careful use of resources. One computer center director complained that his users didn't take advantage of off-peak discounts. On investigation, the cause was found to be that users had quite enough budget to do all their computing at premium rates!

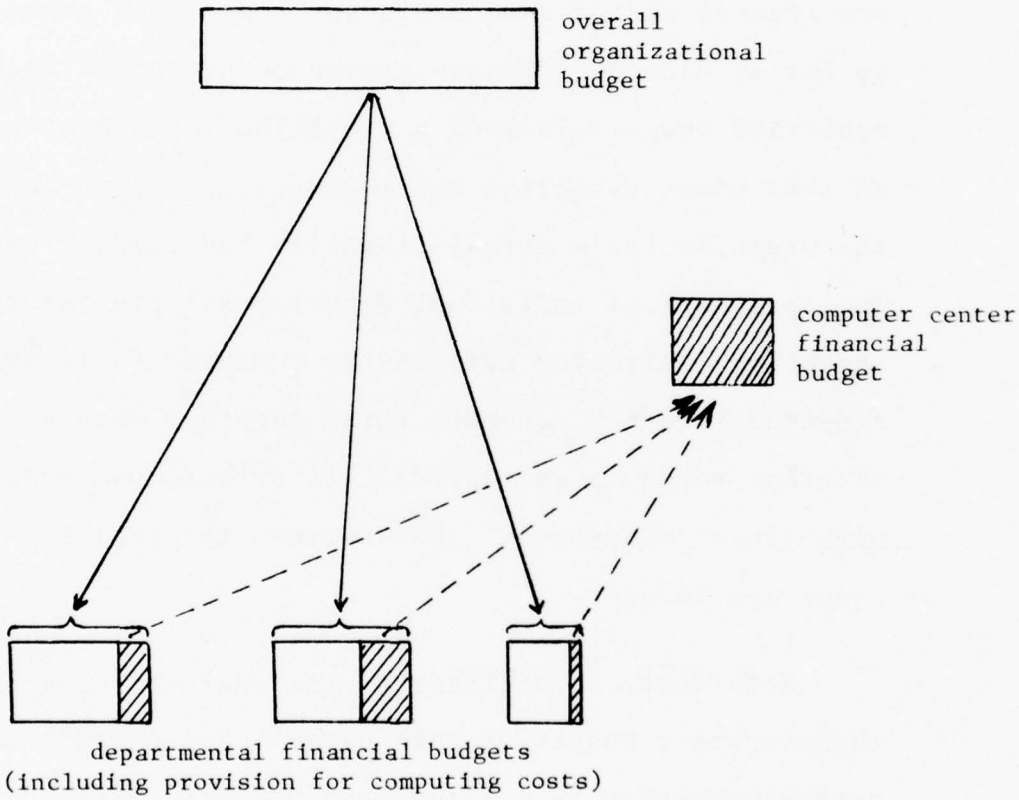
Decentralized Budgeting

The charging mechanism plays a far more central role when computing is budgeted for on a decentralized basis. This approach is depicted -- again in highly simplified form -- in Figure 3.2. Here, internal computer charges are treated on the same basis as 'real' cash expenditures as far as user departments are concerned; departmental operating budgets include a provision for computing costs, so that users establish their computing allocations within the organization's normal financial budgeting framework. The aggregate of individual departments' provisions for computing implicitly defines the computer facility's expected internal revenue; this, together with anticipated external revenues and any deficit or surplus negotiated with senior management, then defines the facility's expenditure budget.

Note that, in contrast to the centralized approach, the computing budget is here established through a bottom-up rather than a top-down process. The basic sequence of steps is reversed: rather than establishing the computer expenditure budget and then from this developing a revenue budget, the latter is established first, and the expenditure budget derived from it.

Figure 3.2

Simplified representation of a decentralized budgeting process for computing



There are important differences between the centralized and decentralized approaches. Under a decentralized approach, user departments have full control over their computing expenditures, within the constraints of their overall operating budgets. While this decentralization of responsibility can bring major benefits, it also relies on users exercising their control effectively. This involves, for example, proper user understanding of the considerations involved in effective use of computing, reporting systems which clearly identify and break down computer costs within users' budgets, and performance evaluation procedures which explicitly focus on managers' effectiveness in using computing.[1]

Computer center management is also greatly affected by the budgeting procedure used. Under a centralized allocation procedure, the computer center's task essentially is to provide the greatest possible quantity and quality of computer services at a pre-determined level of

[1] Further discussion of these considerations will be found in R. L. Nolan: Management Accounting and Control of Data Processing, National Association of Accountants (forthcoming).

funding. When user departments are given direct control of their computing expenditure, the computer center faces an internal 'market' for its services, and thus has a much more entrepreneurial role. Computer center management has considerable autonomy in choosing the scale and type of computing capacity to be provided, but must generate sufficient revenue to support this capacity (taking into account any agreed surplus or deficit); thus the task now entails not just management of costs, but also management of demand. This involves taking account of demand elasticities in setting prices,[1] and may also require external selling of services to make up shortfalls in revenue. Moreover, given the difficulties involved in adjusting computer center capacity and costs from year to year, the center will need to take measures to reduce the uncertainty in future demand. Involving users in long-term

---

[1] Experience at the Massachusetts Institute of Technology has shown that users there typically react to a price change in one of two ways: either they continue at their previous level of expenditure, altering their consumption accordingly, or they maintain their previous consumption, accepting the associated change in expenditure. The MIT Computer Center's revenue appears to divide roughly equally between the two categories, so that a ten percent change in computing rates produces a five percent change in revenues, and a roughly equal opposite change in resource usage.

planning for computing, and encouraging long-term supply contracts with major users can be important tools in this area.

From the viewpoint of senior management, the most important difference between the centralized and decentralized approaches lies in the replacement of a negotiation process to establish the allocation of resources to computing and between computer users, with a market mechanism which establishes this allocation implicitly through the decisions of individual user groups. In theory, given suitable computer pricing, a decentralized budgeting process will produce a cost-effective allocation of resources to computing and between users, working in the manner of Adam Smith's "invisible hand".[1] Thus, again in theory, management should be able to decentralize control over the scale and distribution of computer usage with confidence that the financial budgeting process will produce the 'correct' allocation.

However, the problem is rather more complicated in practice; doubts are often expressed about the ability

---

[1] See S. Smidt: "The Use of Hard and Soft Money Budgets, and Prices to Limit Demand for a Centralized Computer Facility, Proc. AFIPS Fall Joint Comp. Conf., 1968, pp 499-509.

of a decentalized budgeting procedure to produce the best use of resources.[1] First, it is pointed out that management's policy objectives with respect to computing may go beyond a straightforward concern for cost-effectiveness. Furthermore, there is often a feeling -- particularly in non-profit institutions -- that the spending power of individual organizational units does not necessary correspond to the value of their activities to the organization; this feeling is reflected, for example, in the common objection that rich departments or projects will gain at the expense of poor ones under a decentralized budgeting arrangement. In addition, it is sometimes suggested that, because of special characteristics associated with computing, line managers' decisions on computer use will not always reflect its true costs and benefits in relation to other expenditures; this concern appears in statements such as "departmental managers do not appreciate the potential value of computing and need to be encouraged to use it," or "line managers are only concerned in short-term results and will not be interested in computer projects which typically

---

[1] See, for example, the discussion of centralized versus decentralized budgeting in S. Gill & P. A. Samet: "Charging for Computer Time in Universities," Computer Bulletin, Vol. 13, No. 1, (Jan. 1969, pp 14-16.

have long-term payoffs: therefore computing must be controlled centrally."

Since there is no objective way to assess the validity and importance of such objections to decentralized budgeting, its acceptability in practice often depends on how far management feels the results would be consistent with its own view of the correct scale and distribution of computing expenditures. Our own opinion is that, provided adequate measures are taken to help and encourage line management to exercise their control over computing effectively, a decentralized budgeting process will generally produce better allocation decisions than a central computer committee. Moreover, as we have already indicated, decentralization has important side-benefits in bringing the allocation process within the organization's normal budgetary framework (thereby eliminating the extra political tensions inherent in a special computer allocation process), encouraging line managers to involve themselves more closely in their departments' computing activities, and greatly increasing the computer center's incentive to remain efficient and responsive to user needs. It is true that a decentralized allocation process has its imperfections; however, as Goldstein, Jensen and Smith point out in discussing a proposed change from centralized to

decentralized computer budgeting at the University of Rochester, the real issue is the importance of these imperfections relative to all the disadvantages associated with a centralized approach.[1] It should also be pointed out that decentralization of the budgeting process still allows management to exercise some degree of central control where it feels this to be necessary, both through manipulation of the pricing scheme (e.g., subsidizing certain types of computer use) and through direct restrictions of various kinds.

Hybrid Approaches

The centralized and decentralized budgeting arrangements described above really represent the two extremes of a spectrum of possibilities; there are a number of ways an organization can compromise between the two approaches by combining elements of both. The examples below indicate some of these 'hybrid' approaches:

- Budgeting is carried out on a decentralized basis, but central control is maintained over the computer center budget (i.e., the center is

[1] D. Goldstein, M. C. Jensen & D. Smith: Report of the President's Committee on Computing Problems and Opportunities", University of Rochester, 1973, pp. 1-10.

not free to use excess revenue to increase capacity).

- Computing budgets are allocated centrally, but are shown as a line item in departmental operating budgets. This ensures that computer costs are reflected in departmental performance figures and other management information, encouraging line managers to remain aware of these costs.

- A basic allocation of computer dollars is made centrally, but this allocation is deliberately kept well short of users' full requirements. Users supplement their allocations as desired with 'real' money from their general budgets. The aim here is to provide the computer center with some stability of income while still obtaining the benefits of users' trading-off of computing costs and benefits at the margin.

- Computing resources are budgeted on a decentralized 'real money' basis at the level of major organizational units (e.g., corporate divisions or schools of a university), but allocated on a funny-money basis from this level downwards. The assumption

here is that budgetary decisions on computing are best made at the intermediate level.

- Real money budgets and charges are used for some classes of use and/or user, and funny-money budgets and charges for others. An obvious example of this approach occurs in universities, where instructional computing is often funded through funny-money allocations, while computing for funded research projects is charged in 'real money' terms against financial budgets. (A difficulty with this type of arrangement is preventing the use of funny-money allocations for purposes that are supposed to be paid for with real money).

Treatment of Non-Central Resources

Our discussion so far has dealt with one dimension of decentralization in the budgeting process -- that is, the extent to which users are free to choose the level of their computing expenditures. Another important dimension is users' degree of freedom to direct these expenditures as they wish -- for example, toward their own mini-computer installation or purchase of outside computer services -- rather than having to use the central facility exclusively. It should be emphasized that these are largely

independent dimensions. Although funny-money allocation schemes typically cover only the services of the central facility, there is no reason why they should not be extended to include all forms of computer services. Similarly, while charging for computing in terms of 'real money' greatly increases users' pressures to be allowed to purchase their computer services from what they regard as the most cost-effective source, such freedom does not necessarily have to be granted.

The key issue in deciding on the treatment on non-central resources is the same as that discussed earlier: i.e., what arrangement is likely to produce an effective allocation of resources -- in this case between central, distributed, and external sources of computing. Under a funny-money allocation scheme that covers only central services, users have little incentive to seek more cost-effective sources than the computer center (since the former must be paid for from operating funds -- if this is even permitted -- while the latter are effectively free); thus the organization is unlikely to take full advantage of alternatives such as mini-computers and resource-sharing networks. On the other hand, if users are given freedom to choose where they spend their computer allocations, the question arises how far their choices, based on the

relative costs to them of central and non-central sources, will be consistent with the organization's overall interests. The same question obviously also arises under a decentralized 'real money' budgeting arrangement.

As will be shown in the next chapter, the computer center's pricing scheme can to some extent be designed so as to encourage users to make the desired choice between central and non-central services. Nevertheless, many organizations are unwilling to rely completely on the market mechanism, and therefore restrict users in their choice of source. This may be done, for example, by requiring central approval of any computer-related expenditures other than for central services. Alternatively even stronger control may be exercised by giving the central computer department responsibility for supplying all forms of computing resources, including mini-computers and outside services, and requiring all such resources to be obtained through that department.

As with the previously discussed issue of centralized versus decentralized budgeting, management must decide the degree of freedom it will allow users over their source of supply in relation to its overall policy goals for computing and its general management control philosophy.

We would point out, however, that whatever the potential dangers of misallocation associated with allowing users freedom in this area, it does bring with it the important benefits associated with decentralization generally: elimination of a cumbersome and conflict-inducing central control process, increase of line management's accountability for all aspects of their departments' activities, and reinforcement of the market pressures acting on the computer center.

Financial Control Options

Cost-based versus Profit-based Control

There are two main approaches available for controlling the computer facility. Where it operates largely independently of outside competition (i.e., serving no external users and internally facing little or no competition from outside suppliers), control must be based on the facility's detailed cost and revenue budgets and on its performance against these budgets. Alternatively, where the facility competes on a significant scale against outside suppliers for internal and/or external business, it can be controlled as an independent profit center or subsidiary.

Under the first approach the facility is typically set up as an internal cost center.[1] In many early charging schemes, costs were allocated retrospectively based on each previous month's actual costs and usage; however this approach had serious drawbacks from the viewpoint of management control: charging rates fluctuated from month to month with the overall volume of usage, severely weakening users' ability to budget for and control their charges effectively. Therefore charging rates are now normally set at the beginning of each year on the basis of planned costs and resource usage[2]; this provides users with a predictable, stable basis for planning and control. Moreover, under the latter arrangement any deviations from

---

[1] One sometimes finds the computer center set up as a profit center even though it is not subject to the constraints of external competition. In our view, such an arrangement has little point: the center's management cannot safely be given the freedom normally associated with a profit center, since it is a monopoly position with respect to its users; if, on the other hand, its profit-seeking behavior is constrained -- for example by imposing a break-even profit target -- this effectively puts it in the same position as a cost center. The profit center arrangement is sometimes adopted on the grounds that this allows market-oriented pricing techniques such as off-peak discounts, long-term bulk supply rates, etc. However, there is no reason why the same techniques cannot be applied within a cost center framework.

[2] In other words, charges are based on standard rather than actual costs.

the plans used as a basis for rate-setting will appear as an under or over-recovery of costs by the computer center. Any such deficit or surplus can be analyzed to determine the extent to which it is due to volume, efficiency, or cost variances; such analysis provides a powerful tool for management control of the center. Where a component of the deficit or surplus can be directly attributed to a particular user department, it may be appropriate to assign it to that user's budget; the remaining deficit or surplus may be handled by charging it to general overhead or by carrying it forward to be recovered in the next period.[1]

The second alternative, in which the facility is set up as a profit center or subsidiary that competes against other suppliers, is growing increasingly popular -- particularly with large corporations. Competition provides

---

[1] Many computer centers handle variances by allocating them among users on some pro-rata basis, or by making mid-year rate adjustments designed to correct any deficit or surplus. Such corrective efforts seem to have little value, since they will not eliminate the cause of the original variance. Moreover, they have the disadvantage that they create uncertainty for users, weakening the basis for his budgetary control. Serious problems could be caused if a retrospective levy to cover a deficit caused a user's budget to be exceeded in a situation where the budget was legally binding.

a strong spur to high quality performance at the computer center. Moreover, the arrangement provides management with a direct measure of the center's performance -- namely its profitability.[1] On the other hand, the profit center arrangement encourages the computer facility to act in its own interests rather than those of the organization as a whole: there is a danger that this will result in less effective computing if users are not sufficiently sophisticated to play effectively the role of 'consumer' in an arm's-length buyer-seller relationship with the center.

An important issue under both the cost center and profit center approaches is the basis used to establish the costs which the computer facility is required to recover (or earn a profit on). Since the facility's rates

---

[1] In some organizations where the computer center does not actively compete with outside suppliers, the test of the market is applied by directly setting charging rages at a level corresponding to external market prices (the latter being assessed through a survey of comparable commercial services). This again allows control of the center on the basis of profitability. The arrangement does, however, depend on the existence of commercial services which are directly comparable with those offered by the center.

are based on these costs, decisions such as whether to include organizational overhead charges or hardware depreciation costs can significantly affect the level of users' charges. In a real money budgeting environment, this will in turn influence the extent to which users utilize the facility's resources.

One consideration in deciding the computer center's cost recovery target is the effect on the cost to users of computing relative to other types of activity, and therefore on the balance struck between these alternatives. More significant, however, is the effect on the competitiveness of the computer center's services relative to alternative sources of computing. This consideration may be important even when users' freedom to select alternative sources is restricted, since rates which users regard as uncompetitive will still generate dissatisfaction with the center and strong pressures for permission to go elsewhere.

The general question of pricing central computer services in competition with non-central alternatives is discussed in Chapter 4. For the moment we will mention only one point which arises in this context -- namely, that it is important that the basis used in establishing the computer center's costs is comparable with that used in

representing the costs of alternatives. For example, a proportion of general overhead is often allocated to the computer center, so that the rates it charges users reflect this allocation of overhead as well as the direct costs of the center. On the other hand, alternatives such as a local mini-computer installation or use of a service bureau may be represented in users' budgets at their basic direct cost. Such a discrepancy in treatment will bias users against employing the central services even when these would in practice be more cost-effective. In one large corporation, the corporate data center instituted a charging scheme which involved recovery of a share of corporate overhead in addition to its direct costs. The divisional data centers' cost recovery targets, on the other hand, did not include any allocation of corporate overhead. The resulting rate differentials caused a major shift of usage away from the corporate center, even though the true costs to the organization of processing at this center were no higher than for the divisional centers.

Wholesale-retail Schemes[1]

Wholesale-retail schemes are a recent innovation of considerable potential value to organizations with large-scale, diverse computing activities. Under such a scheme, the functions associated with production of computer services (i.e., running the data center) are separated from those associated with delivering these services to end users (i.e., providing user support, charging, etc.). The computer center (or centers) no longer deals directly with users; instead it supplies its services on a bulk-supply or 'wholesale' basis to distributor organizations, which in turn 'retail' these services to their local users. The 'wholesaler' and 'retailer' organizations are financially independent, dealing with each other on an arm's length contractual basis.

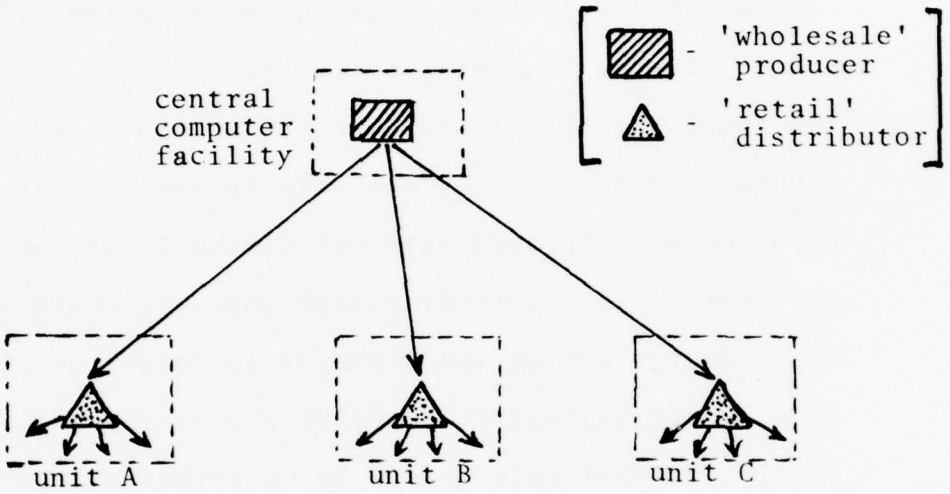
Figure 3.3(i) illustrates the way such an arrangement can be used within an organization. Here a central computer

[1] Further discussion on this topic will be found in D.L. Grobstein and R. P. Uhlig: "A Wholesale Retail Concept for Computer Network Management," Proc. AFIPS Fall Joint Computer Conf., 1972, pp. 889-898, and in a paper by E. Stefferud presented at the ACM SIGMETRICS Technical Meeting on Pricing Computer Services, Nov. 20-21, 1975, and published in Performance Evaluation Review, Vol. 5C, No. 1, March 1976, pp. 31-70.

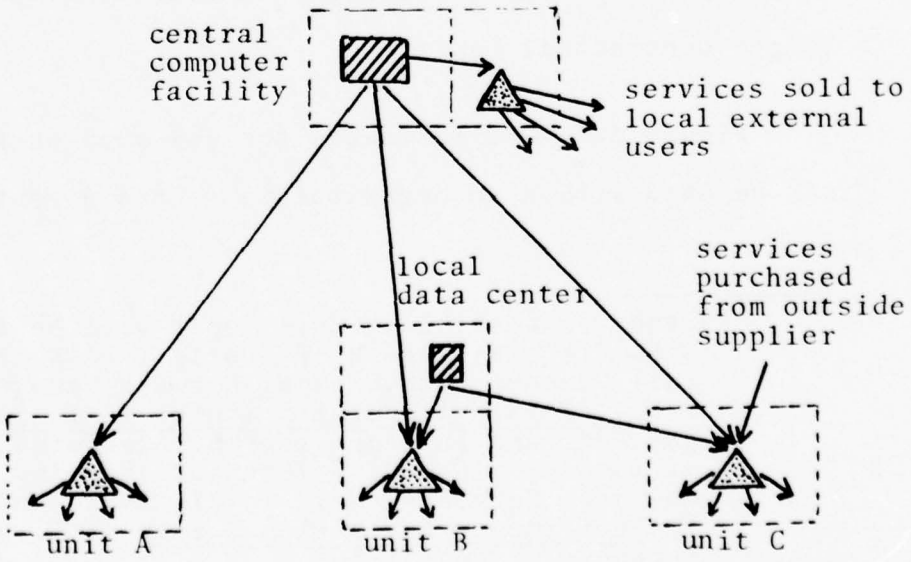
Figure 3.3

Illustrative wholesale-retail schemes for control of computing

(i)



(ii)



facility supplies services in bulk to three organizational units. (Within a corporation these units might be divisions or functional units; within a university they might represent, for example, two schools of the university plus the administrative branch, or alternatively three campuses in a multi-campus institution). Each organizational unit buys an overall volume of computer services from the central facility, and itself controls the way these services are distributed. Thus each unit can tailor its handling of computing to its own particular needs and policies. In particular, each unit can make its own decisions on pricing, budgeting, and financial control of computing. For example, one unit might decide to provide computing free, a second to use a funny-money allocation scheme, and a third to budget and charge for its computing internally in real money terms; one unit might decide to subsidize its computing to users, while another insists on recovering the full cost of providing these services; different units might also employ quite different bases for charging.

The advantages of distributing control in this way become particularly strong in more complex situations such as that shown in Figure 3.3(ii), where there are several

computer facilities, each shared between several units, and where the organization is also involved in supplying services to, and purchasing them from, other organizations. Through the wholesale-retail scheme, issues relating to control of the computer centers, supply to external users, and allocation and use within the three user areas, can be effectively separated. Moreover, each user area is able to handle its computing activities through a single uniform pricing and budgeting framework, even though it obtains computer services from several different sources. This greatly simplifies what could otherwise be an extremely complex management problem.

Chapter 4

DESIGN OF A PRICING SCHEME

We now turn our attention to the central component of any charge-out system -- the pricing scheme through which users' charges are calculated. We begin by outlining some basic principles underlying the design of a pricing scheme. The remainder of the chapter discusses the development of a basic resource-related price structure, alternative pricing methods, pricing of service priority, charging for systems and programming resources, and finally some special pricing problems.

Underlying Principles

Objectives in Pricing

It was pointed out in Chapter 1 that for the purposes of effective design, charging is best regarded as a control mechanism rather than simply as a cost allocation process. If charging is viewed as a control mechanism, it becomes clear that the key consideration in deciding the basis for calculating charges, and the level of charging rates, should be the likely effect on users' decisions. The central objective must be to establish a price structure which,

through its influence on users, will provide maximum assistance in producing the desired pattern of computer use.

The pricing scheme can influence users at several levels, including such areas as:

- deciding whether to implement a new computer application, or choosing between alternative applications.
- choosing between use of the central installation or an outside resource.
- balancing the use of computing and manual effort (for example, in a statistical analysis problem).
- deciding between technical strategies which differ in the mix of resources required (as, for example, in 'fine-tuning' a program's structure to optimize the balance between core size and input-output activity, or in choosing between a batch or an on-line approach to data input).
- deciding what priority to request when submitting a job, or what time to start a timesharing session.

In all such areas, an ideal pricing scheme would always motivate users to act in a manner consistent with manage-

ment's objectives.[1] This ideal obviously cannot be achieved in practice. Nevertheless, in designing a pricing scheme its effect on user behavior needs to be considered explicitly; this will maximize the effectiveness of charging, and minimize the need to supplement the charge-out system with direct controls and restrictions on users. As will be clear from subsequent discussion, such an approach requires consciously pricing, rather than simply costing, services and may involve using many of the techniques employed by commercial service bureaux.

Relation of Prices to Costs

On the whole, computing resources will be used most cost-effectively if users' decisions are based on a price structure which accurately reflects costs. For example, a user's decision on whether to use the computer in a particular application will normally be based on his estimates

[1] Clearly users will have discretion in all the decision areas described above only where control over computing is highly decentralized. When the range of decisions delegated to users is more limited -- as, for example, within a funny-money budgeting scheme -- the pricing problem is correspondingly simplified. Many of our observations on pricing in this chapter may not be relevant where the role of charging is comparatively restricted.

of the costs and benefits involved. Since the 'costs' are represented to the user in the form of his computing charges, these charges ideally should reflect the actual costs to the organization of the resources involved. Similarly, a user's choice between technical strategies involving different mixes of resource usage would normally best be based on rates reflecting the relative cost of the various resources involved. It is clear, therefore, that the computer center's cost structure must be used as the basis for establishing prices.

This does not, however, imply that the price structure must necessarily always reflect costs. There may be instances where deviations from a strict cost basis for pricing will lead to better use of resources, or will better support policy objectives. Thus pricing requires a flexcible approach: strict considerations of cost need to be supplemented with subjective judgements as to what is most likely to yield the desired results.

The examples below illustrate some of the circumstances in which deviations from cost-based pricing may be justified. Further examples will appear throughout the chapter:

- Deviations from cost may be advisable where computing policy does not follow strictly economic criteria. Consider, for example, the case of a small college that decides, for prestige or other reasons, to maintain its own computer center even though its small scale implies a higher cost than the alternative of using outside resources. Under these circumstances, it may be advisable to subsidize the computer center so as to bring internal charging rates down at least to a level comparable to outside rates; otherwise, users will continually complain of excessive internal charges, and will whenever possible use other sources, to the detriment of the organization as a whole. The subsidy payment here explicitly represents the costs incurred by management's policy of maintaining an in-house computing capability.

  A further example is that of a university which wishes, for reasons of institutional policy, to boost the use of computing for instructional purposes. Special rates for instructional use, subsidized from a central fund, could be used to support this objective.

- Special pricing arrangements can be used to encourage users to take actions that will benefit the organization as a whole. For example, it has already been pointed out that long-term commitments from major users can greatly reduce the risk of the computer center being left with excess capacity through unanticipated fluctuations in demand. Discounts can be employed to induce users to make such commitments, in effect transferring the risk from the computer center to themselves.

  Another example arises in experimentation with new technology or pioneering 'state-of-the-art' applications. The experience gained from such activities potentially benefits the whole organization, but involves the initial users in considerable risk. It would be appropriate for the organization to promote experimentation by subsidizing these activities.

- Prices may be manipulated to regulate demand. It was pointed out in Chapter 1 that, because computer capacity cannot be changed easily except in occasional, large increments, it may sometimes be necessary to accept a deficit or surplus at the

computer center (implying prices which reflect less or more than full cost), in order to match demand to the available capacity. An example is provided by Stanford University, which at one point quadrupled the capacity of its computer center to allow for long-term growth. Because the new configuration was initially heavily under-utilized, charges based on an apportionment of full cost would have been very high and would have discouraged demand in a situation calling for rapid growth in volume. Instead, Stanford deliberately kept its rates down, even though this meant subsidizing the center while usage expanded.

Pricing may also be used to regulate demand for a particular resource within an installation. For example, a situation analagous to the Stanford experience described above can arise with shared software and data bases whose usage is initially low, but is expected to grow with time. Where such resources are explicitly charged for, it would be reasonable to set prices at the level expected to prevail in the long-term when usage has expanded, even though at the low level

of initial usage, costs would not be recovered.[1]

As another example -- this time dealing with insufficient rather than excess capacity -- an installation might face a temporary shortage of disc storage space while waiting for demand to grow sufficiently to justify installing another storage unit. During this period, it may be necessary to raise disc storage charges well above their accounting costs (recognizing the temporary scarcity value of the resource) in order to ration the available capacity.

Pricing may also be used to shift demand from one resource to another. It is common for one resource to be a bottleneck in an installation while others are relatively under-utilized: for example, throughput may be limited by input-output capacity, leaving the remainder of the installation resource -- core, central processor, etc., -- relatively idle. This situation can be

[1] If a large deficit is unacceptable to the organization, it can be amortized over several future periods, and if necessary, recovered in future charges.

corrected by weighting the price structure heavily toward input-output, thereby encouraging users to reduce their usage of this resource -- for example, by shifting tables, arrays, etc., from disc to main memory.

- Pricing of services to external users may have to take account of competition as well as costs. The organization's objective with respect to external users of its computing resources may be to maximize the financial contribution provided by these users rather than simply to charge them an equitable proportion of costs. In this case the price structure for these users will need to consider the type and volume of business the center wishes to attract from outside and the market rates for this business, as well as internal costs.

Characteristics Required in a Pricing Scheme

It is important that users accept the basis for their charges, and actively use them as a tool for controlling their computing activities: sophisticated pricing will be of little value if charges are simply ignored. These considerations point to several necessary characteristics in the design of a pricing scheme.

First, users must generally accept the pricing scheme as equitable. This does not necessarily imply that prices must adhere rigidly to costs, but any deviations from a cost basis -- particularly involving discrimination between users -- will probably have to be justifiable to the user community. Obvious inequities that result in a class of users feeling it is bearing more than its fair share of costs, are likely to produce strong reactions and prove a continual source of tensions.

More significantly, the scheme must be understandable to users if they are to be able to trace the cause of variations in charges, intelligently seek ways of reducing them, and generally effectively budget for and control their computing activities. As we have already said, pricing schemes are very frequently made overly complex and technical in an attempt to measure accurately every aspect of a task's use of system resources, when such accuracy may be unnecessary, and positively harmful if achieved at the expense of user comprehension. A recent survey of industrial firms revealed that:

> "The majority of user/managers interviewed did not understand their data processing chargeout bills anywhere near the level assumed by the designers of charge-out systems. Consequently, the actual control effects of most charge-out systems were a sham. User/manager control actions were assumed to be

> motivated by deliberation on charge-out bills; in fact, the actions were usually taken for other reasons."[1]

The importance of a concern for user understanding cannot be overemphasized.[2]

Another important characteristic of a pricing scheme is that it should, as far as possible, yield charges which vary only with factors that are controllable by the user.[3] Thus, if a user makes his program more efficient, or eliminates an unnecessary report from a management information system, his charges should change to reflect his actions -- and should do so in a predictable manner. On the other hand, changes which are outside the user's control, such as variations in the installation loading or

---

[1] R. L. Nolan: Management Accounting and Control of Data Processing, National Association of Accountants (forthcoming)

[2] User understanding involves more than suitable design of the pricing scheme. Educational efforts will be needed to explain the scheme and introduce changes. It may also be helpful if computer center staff work with line managers in reviewing their charges, and budgeting for future expenditures.

[3] This is, of course, a principle which applies to any budgetary control system.

configuration changes, should affect the charges for a given processing task as little as possible.[1] Variations in charges over which the user has no control lead to frustration and seriously hinder effective budgetary control.

A particular requirement that arises from the controlability consideration is that the pricing scheme should give reproducible results. That is, when a particular job is run on the installation under varying conditions, the charge for the job should remain constant as far as possible. We discuss in the next section some of the difficulties involved in satisfying this requirement within a multiprogramming environment.

Finally, closely related to controllability is a need for stability in the charging structure, particularly within a budgetary period. Price adjustments are needed to take account of changes in the cost structure (e.g., due to configuration changes) and shifts in the pattern of demand. Moreover, controlled experimentation with pricing can greatly increase the computer center's understanding of

---

[1] Note that his consideration may lead to a basis for pricing that does not strictly reflect the cost of resources used. For example, if a job incurs greater system overheads when the machine is busy, this should not be reflected in increased charges.

its market. However, frequent change in the charging structure weakens users' ability to budget for and control their charges, and frustrates those who base design decisions on a particular charging structure, only to find it change to favor a different technical approach a few months later. Therefore the charging structure should be altered as infrequently as possible -- particularly in the middle of a budgetary period -- and maximum notice should be given when a change is necessary.

Establishing a Price Structure

In any installation there is a large range of resources that could, in theory, be charged for individually. In practice, given the difficulty of accounting for the use of each resource, and the need to keep the charging structure simple, installations typically select a subject of key resources whose usage forms the basis for billing.[1] Associated with each resource is one or more utilization measures (some common examples are shown in Figure 4.1); these measures are combined with associated unit prices to define the formula, or 'algorithm', through which charges are calculated.

---

[1] The alternative of charging in terms of application-related parameters rather than resources is discussed in the next section.

Figure 4.1

Common resource utilization measures

Resource category	Measure	Definition
Central system	CPU time	Central processor time used
	I/O time	Total time job occupied input-output channels
	I/O requests	Total number of input-output requests issued by job
	Program size	Job's memory requirement
	Memory time	Integral over time of job's memory requirement
	Productive memory time	As above, but limited to periods during which job had control of CPU
	Swap time	Total time spent swapping pages for job (in virtual memory system)
Unit record peripherals	Cards read	)
	Cards punched	)
	Lines printed	) (self-explanatory)
	Plotting time	)
	Plotting units	)
Data storage	Block-days	Integral over time of number of blocks of on-line disc storage occupied by user
	Peak blocks	Maximum number of blocks of on-line storage occupied during the accounting period
	Tape-days	Integral over time of number of tape reels reserved by user
Terminals and ports	Connect time	Total time user logged onto system from an on-line terminal
Operator time	Set-ups	Number of special requests to mount tapes, disks, special printer stationery, etc.

We describe below how a basic price structure of this sort can be developed, focusing on the considerations involved in selecting the resources to be charged for, designing suitable utilization measures, and setting the prices associated with these measures. We are limited here to a fairly brief discussion of what is a highly complex technical problem: there is an extensive literature available as a source of further information.[1]

---

[1] See, for example:

N. R. Nielsen: "Flexible Pricing: an Approach to the Allocation of Computer Resources," Proc AFIPS Fall Joint Computer Conf., 1968, pp. 521-531;

J. T. Hootman: "The Pricing Dilemma," Datamation, August '69, pp. 61-66;

C. B. Kreitzberg & J. H. Webb: "An Approach to Job Pricing in Multi-programming Environment, "Proc. AFIPS Fall Joint Computer Conference, 1972, pp. 115-122;

H. M. Gladney, D. L. Johnson & R. L. Stone: "Computer Installation Accounting," IBM Systems Journal, No. 4, 1975, pp. 314-339;

M. M. Lehman: "Computer Usage Control," Computer Journal, Vol. 16, No. 2, (May 1973), pp. 106-110;

G. K. Wiorkowski & J. J. Wiorkowski: "A Cost Allocation Model," Datamation, Aug. 1973, pp. 60-65.

Selecting Resources to be Charged For

The key consideration in deciding what resources are to be charged for directly is that it is on this subset of resources that users' attention will be focused by the charge-out system. Thus the resources chosen should reflect the key cost elements in the installation's budget. The main emphasis should be on the most heavily utilized resources -- i.e., those that impose the main limitation on installation usage. For example, in some dedicated time-sharing installations, capacity is primarily limited in terms of a maximum number of terminals which can be supported at any one time, irrespective of the type of work each is doing. In such a case charges are best made purely on the basis of terminal connect time, without reference to processor, core, and I/O channel utilization.

Because hardware used to be the main cost element in computer installations, charging has traditionally been based entirely on machine usage, with the cost of services such as application software libraries and user support typically 'bundled' into the machine usage charges. However, the cost of these non-hardware-related services has grown to the point where it is as significant as the hardware in many installations' budgets. There are therefore now strong arguments for charging directly for

these services, to control the demand for them, and to avoid subsidy of those who use them by those who need only basic hardware resources.

Designing Utilization Measures

Associated with the selection of resources to be charged for is the choice of parameters to measure the usage of these resources: for example, use of a machine's input-output channel capacity can be measured in terms of number of I/O requests issued, number of characters transferred, total time over which I/O channels were tied up, etc. Several considerations arise in designing suitable utilization measures.

First, it must be borne in mind that, just as the choice of resources charged for affects the direction of users' attention, so does the choice of utilization measures: for example, a user charged for I/O in terms of total characters transferred is unlikely to be concerned with choosing efficient blocksizes so as to minimize the number of separate I/O requests required for transfers. A particular point that arises here is that where users are able to prevent others from using a resource even when they are not using it themselves, they should pay a

penalty for this. This principle, which Hootman[1] refers to as the "principle of demurrage", would for instance imply that where a user can reserve a printer for use during processing, and might leave it reserved over a full processing sequence when he could have released it at an earlier stage, then he should be charged on the basis of the time the printer was reserved rather than simply in terms of the more conventional measure of total lines printed.

The issues of understandability, controllability, etc., discussed in the previous section must also be considered. For example, number of lines printed is probably a more meaningful measure to the user than total print time. Moreover the number of lines printed is much more directly controllable by the user: print time is dependent on factors such as operator efficiency, and if the installation has several printers of different speeds, print charges will vary depending on which is used. A particularly difficult problem is the identification of measures that accurately reflect resource usage while yielding reproducible results. The reason for this is that in a multi-programming environment, the time taken to run a job and the resources it uses (particularly if system overheads such as page

[1] Op. cit.

swapping are taken into account), are directly dependent on the installation's overall workload at that time. The problem can be illustrated with reference to the measurement of a job's use of main memory -- an important limiting resource in most installations. The measure best reflecting the use of this resource would be the job's program size multiplied by its total run-time. However, charges made on this basis would vary from run to run since the job's run-time changes with the installation's loading. Therefore installations generally multiply the job's program size by its processing time rather than overall run-time, thereby excluding periods that the job is inactive while other jobs have control of the machine.

A further factor that limits the choice of utilization measures is that the requisite data must be relatively easily obtainable from statistics provided by the operating system, or from other sources such as a hardware monitor, operator console logs, etc. As an example, the System Management Facility (SMF) of IBM's Operating System 1370 (OS) provides a large range of statistics that are generally used as the main input to charging routines for IBM installations. However, the only statistic available from SMF relative to input-output activity is a count of I/O requests (EXCP's) for each task: thus this particular

measure is almost always used, irrespective of the merits of the theoretical alternatives.

A final point to be made is that the same set of utilization measures need not necessarily be used for all the types of work carried out by an installation. For example, quite different charging formulae might be used for batch and time-sharing work. The charging system described by Lehman[1] embodies four separate schemes for the four main categories of use within the installation.

Establishing Prices

Once the basis of charging has been defined the process of establishing prices is basically that involved in the calculation of standard costs for any service center. The installation's overall revenue requirement (i.e., its costs, adjusted by any planned deficit or surplus) is allocated among the resources to be charged for to give a revenue share for each element in the charging formula. Dividing each share by the anticipated utilization of that resource yields the associated price per unit.

In the allocation process some costs will be clearly attributable to a particular resource, while others

[1] Op. cit.

(including any general subsidy or profit) will have to be distributed across the various charging elements in some way. As was pointed out in the previous section, the allocation process needs to be carried out flexibly, taking into account not only the general desire that the price of a resource should accurately reflect its cost, but also the overriding consideration that the price structure should induce the desired behaviour among users.

Alternative Pricing Methods

There are other ways of pricing an installation's services that can be used instead of, or in conjunction with, the approach described above. These alternative methods can have considerable advantages in dealing with certain types of resource or service.

Output-related Pricing

Pricing services in terms of physical resource usage is most appropriate where the installation is primarily supplying raw computing resources, with users controlling the way these resources are employed. An alternative, where the installation provides information processing resources rather than raw resources (as, for example, with administrative EDP applications), is to price services in terms of the outputs provided rather than the resources

used in producing these outputs. For example, a routine accounting application might be charged for on the basis of the number of accounting transactions processed, with the set price per transaction covering validation, update, and production of standard reports. Similarly, payroll processing might be priced at a set fee per employee, with additional charges to cover non-standard inputs, such as special deductions, that create additional work. Again, information retrieval services might be priced in terms of a charge per inquiry, and/or per item retrieved.

The key advantage of this approach is that users' charges are based on parameters that they can understand and control, rather than on resource-related measures that have no direct meaning in terms of the application involved. To quote a data processing executive within a large insurance company:

> "Our users have been budgeting for premiums collected for about 112 years, and they should be reasonably competent at doing that. They shouldn't have to think in terms of the units that the people in the back rooms use to generate the product."[1]

---

[1] Remark by F. Kirshenbaum, quoted in J. C. Emery and H. L. Morgan: "Management and Economics of Data Base Management Systems", in D. A . Jardine (ed.): Data Base Management Systems, American Elsevier/North Holland, 1974, p. 192.

Moreover with output-related pricing the computer center, which typically controls the development and operation of the application systems, has a financial incentive to ensure that the application work is processed efficiently; when the center charges directly for resources consumed, the user bears the cost of inefficient programs and operational procedures, even though he cannot control these factors.

These advantages have understandably led to the widespread adoption of output-related pricing among commercial bureaux offering EDP services. Its use within internal computer centers is as yet very limited -- probably largely because such centers tend to regard charging primarily as a cost allocation mechanism rather than as a control process. However, research has shown that where output-related pricing has been adopted -- typically within organizations with fairly mature data processing activities -- it has generally produced a major improvement in the acceptability and effectiveness of charging.[1]

---

[1] See R. L. Nolan: Management Accounting and Control of Data Processing, National Association of Accountants, (forthcoming)

Pricing of services in output-related terms involves estimating the resources that will be used for each unit to be charged for (e.g., per transaction, per inquiry, etc.), and then applying standard unit resource costs established for pricing purposes within the computer center. Clearly the units used for charging will have to have fairly well-defined resource requirements to allow reliable estimation of costs. In an information retrieval application, for example, it will probably be necessary to charge partly in terms of number of inquiries and partly in terms of number of items retrieved, to reflect the effect of both these parameters on resource usage. Moreover, if there are different types of inquiry with widely different processing costs, a separate rate may have to be established for each type.

As with resource-related pricing, the charging parameters chosen need to be related to those factors under the user's control that affect the volume of resources used. For example, if punching and processing costs are affected by the accuracy and legibility of coding forms supplied by the user, then it is desirable that charges be based on some measure of this (e.g., number of validation rejections), as well as simple volumes. Similarly,

if late submission of data by users causes scheduling difficulties, some standard scale of surcharges may be applied to discourage this practice.

A final point is that there is no reason why output-related pricing should not be used in parallel with resource-related pricing where the latter method is still more appropriate for some users (as, for example, with a university installation serving both academic and administrative users). Furthermore, it may often be appropriate to employ a combination of output-related and resource-related prices. In an on-line information retrieval application, for example, a terminal connect charge may be needed in addition to the output-related charges, to discourage users from tying up access ports longer than necessary. Similarly, in batch processing applications involving variable amounts of printing, it may be easiest to charge directly for this aspect of resource usage.

Flat-rate Pricing

In many cases it is possible to offer users the opportunity to contract for resources on a flat-rate, usage-independent basis. For example, terminals and access ports, and blocks of disc storage, can be priced on an annual lease basis. Similarly, use of a specialized software

package can be charged for through a fixed 'usage license' fee. It is even possible to price an entire machine in this way: Dartmouth College currently offers the following alternative in its rate schedule:

> "Flat rate contracts are available on an annual basis to educational institutions, a contract provides dedicated access to a port on the system for 9 or 12 months, up to 100,000 words of storage for files (distributed as the institution desires among its user numbers), and a run-time limit of 8 seconds per job. The run time limit, which applies each time RUN is typed, can be up to 64 seconds for two users in each flat-rate group. Background use is not included in the flat-rate package but, like extra storage, may be allowed at an extra charge upon application to Kiewit. The charge for a flat-rate contract for 12 months is $600.00 per month, and for 9 months is $660.00 per month."[1]

This form of pricing can provide significant benefits. It reduces the uncertainty in revenue and demand faced by the computer center, and also the uncertainty in computing charges faced by the user. Moreover it allows the computer center to reward those users who are able and willing to commit to a given level of resource usage (through flat-rate prices that represent a lower cost per unit than the normal usage-related prices), while still providing other users with the flexibility they require (at a correspondingly

---

[1] Kiewit Computation Center, Dartmouth College: rate schedule in effect July 1, 1976

higher cost). From the organization's viewpoint, the fact that users' charges are independent of usage means that they are more likely to fully utilize the resource. Note, though, that this same fact may mean that direct restrictions have to be imposed to limit users' consumption of resources; the Dartmouth example quoted above provides an illustration of this.

Note that it is possible to combine flat-rate with usage-related pricing, in order to gain some degree of stability in computer center revenue while still obtaining the advantages of usage-related pricing in controlling user behavior. For example, charges may be made on a usage-related basis, but with a fixed minimum charge. Or flat-rate charges may be set to cover the fixed costs involved in providing the service, and supplemented by usage-related charges intended to recover the variable costs.

Differential Pricing

Differential pricing refers to any situation in which the same resource or service is priced at different rates under different conditions. This technique can be used for a number of purposes -- e.g.:

- Indirectly charging for certain resources. For example, higher rates may be applied to those users

who require user support services than those who only wish to make use of the installation's processing capacity. Or a surcharge on normal rates may be made for jobs employing a particular software package. These differentials provide a means of unbundling these resources from the normal processing charges, without going to the lengths of separately accounting for their use.

- Discriminating between users (when this is desired). An obvious example is the use of different rates for internal and external users of the center.

- Discriminating between different levels of resource usage. An example is a disc storage charge which reduces as a user's file volume increases (e.g., 10¢ per block per day for the first 1,000 blocks, plus 5¢ per block per day above this level): this provides a discount for large-scale users. Similarly, the opposite technique (i.e., increasing storage rates) might be used to actively discourage large-scale use of disc storage -- perhaps to encourage such users to move their files off-line.

- Distinguishing between different qualities of service, in terms of priority, response time, time of day, etc. (This technique is discussed more fully below).

Pricing Service Priority

The Priority Problem

Any shared computer facility faces the problem of assigning priorities to the work of competing users. The problem can, of course, be ignored by operating on a simple first-come, first-served basis. However, an explicit priority assignment mechanism can improve the overall value of the installation's service by taking into account differences between users' turnaround needs. For example, if one user badly needs his results to continue urgent work while another is relatively indifferent to a delay in obtaining his, it would clearly be in the organization's interest to process the first user's work first, irrespective of who was in the queue earlier.

A related problem occurs with installations that face variations in demand over daily, weekly, or seasonal cycles, or peaks associated with organizational activities such as payroll preparation or end-of-month reporting.

These variations result in congestion -- with its associated costs in terms of users' time and inconvenience -- during peak periods, and/or excessive idle capacity -- with its associated opportunity costs -- at other times. Given these costs it is normally worthwhile shifting work from high-demand to low-demand periods -- for example by deferring work for overnight or weekend processing -- even though this inconveniences the users affected. Again, the overall value of the installation's service will be maximized if deferral is applied selectively to those users who would be least inconvenienced.

Direct vs. Indirect Control

The discrimination between users involved in both the above areas can be controlled entirely by the computer center. For example, the center can establish set criteria as to which categories of work are to be given high priority, which to be deferred overnight, etc. However, this approach has the disadvantage that it places on the center the burden of judging between competing users' claims for priority. A generally preferable alternative -- particularly in a job-shop environment -- is to control the allocation of service priority through the charging mechanism. By offering a range of service qualities, each associated with a different charging rate, users can be

left to decide for themselves what service level they feel is worth paying for given their particular needs. Decentralization of control in this way provides the advantages that have already been mentioned relative to decentralization in other areas -- i.e., replacement of an unwieldy, often inflexible, and highly political direct control process with an impersonal mechanism which influences users through economic incentives while leaving control in their hands.

Pricing Techniques and Considerations

Several methods of pricing service priority are available. They include:

- multiple batch processing input queues that are given different priorities in scheduling, processing rates rising with priority. Users specify the priority when submitting a job and pay the appropriate rates.

- reduced rates for jobs submitted for evening, overnight, or weekend processing, and/or premium rates for peak-period processing.

- varying rates for time-sharing sessions conducted at different times of day or week, to encourage users to login during off-peak periods.

In each case, the price differentials between service classes must be sufficiently high to eliminate congestion in the high-priority queues and peak periods, so that high-quality service is available at least in these premium classes. More generally, price differentials need to be set so that the range of service levels available satisfies the full spectrum of needs within the user community. Experience has shown that large differentials -- up to 60% discounts and 250% premiums -- may be necessary to produce an adequate spread. Clearly the overall rate structure must also satisfy the requirement that, at the anticipated utilization of the various service classes, the installation's revenue requirements will be met.

A further consideration is that tasks that place a heavy load on system resources (e.g., those with large processing, core, or input-output requirements) have a disproportionately large impact on job turnaround times (or response times in the case of time-sharing use). It is therefore worthwhile making specific efforts to shift these particular tasks out of high-priority and peak-period service classes. This can be done through direct restrictions (e.g., "no priority 1 jobs over 5 minutes' run-time", or "all job output over 100 pages will be stored for overnight printing.") Alternatively the price mechanism

can be used to achieve the same result in a more flexible manner. For example, Nielsen describes a scheme used at Stanford University to discourage long jobs during the busy daytime hours:

> "Rather than charging one of a set of rates for the entire amount of compute time, a step function is used. There is a base rate for daytime computing and a somewhat lower base rate for overnight computing. During the day the base rate is increased by 50% for all time in excess of five minutes and by 100% for all time in excess of ten minutes. During the night the rate is reduced by approximately 7% of the base rate after every thirty minutes of continuous processing by a job. Thus long jobs, which are discouraged during the day, are encouraged at night."[1]

Charging for Systems & Programming Resources

So far our discussion has centered around the problems of charging for processing resources -- hardware, common software, and associated support services. We now turn to the considerations involved in charging for resources used in developing and maintaining application systems, where an installation provides these services.

When to Charge

Explicit charging of system development and enhancement costs forces users to weigh these costs (along with

[1] N.R. Nielsen: "Flexible Pricing: An Approach to the Allocation of Computer Resources", Proc. AFIPS Fall Joint Comp. Conf. 1968, pp 521-531

operating costs) against the expected benefits in deciding whether to commission any work. Charging of these costs is therefore generally advisable in the interests of ensuring effective use of systems and programming resources. However, the effectiveness of such charging is dependent on users having a reliable estimate of implementation costs and the ability to evaluate the anticipated benefits.[1] Where these conditions are not satisfied -- as, for example, in an organization new to the use of computers, with users unfamiliar with computer applications -- it may be preferable to centralize decisions on the deployment of systems and programming resources, in which case the associated costs are normally best treated as part of the organization's general overhead. Also, where systems and programming resources are only used on a small scale, and constitute only a small proportion of total computing costs, it may not be worthwhile charging for them explicitly. In this case the associated costs can be bundled in with the processing charges as is typically done with user support services.

---

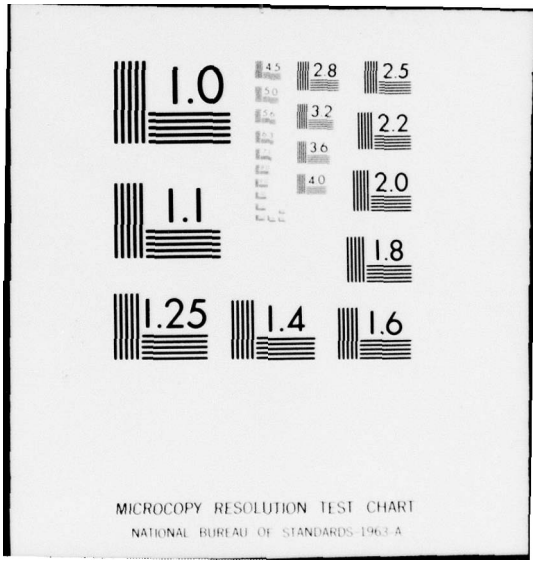
[1] Ideally some form of cost-benefit analysis should be a formal requirement prior to the initiation of any major development project.

AD-A038 082 WHARTON SCHOOL OF FINANCE AND COMMERCE PHILADELPHIA P--ETC F/G 9/2
CHARGING FOR COMPUTER SERVICES.(U)
FEB 77 D BERNARD, J C EMERY, R L NOLAN N00014-75-C-0462
UNCLASSIFIED 77-02-06 NL

2 OF 2
ADA038082



END
DATE FILMED
4-77


1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

Different considerations apply to the treatment of ongoing maintenance costs for application systems. Given that these systems are normally developed by the computer center, there are strong arguments for maintenance costs being borne by the center, since they depend mainly on factors under its control (i.e., the quality of the original programming and testing). In this case maintenance costs are treated as an overhead within the computer center. Alternatively a fixed maintenance charge may be made for each system, with the computer center bearing any variance between this charge and actual costs.

Basis for Charging

Returning to the treatment of development costs, these are usually charged out in terms of man-days of systems and programming resources used, along with non-personnel-related costs such as computer time used for programming and testing. The daily rates for systems and programming staff should be based on expected costs, including overheads such as fringe benefits, office space, etc.; the rates should also allow for a proportion of unbillable time due to lack of work, vacation, sickness, training, etc. Standard rates should be set for each grade of staff (e.g., junior programmer, senior programmer, etc.), rather than using a single average figure for all skill levels or establishing

a separate salary-related rate for each individual. (The problem with the latter approach being that project costs then depend on personnel assignment decisions over which neither the user nor the project manager has any control).

As an alternative to charging on a 'time and materials' basis, the computer department may quote a fixed implementation price at the beginning of the project, and bear any variance between this and the actual implementation costs. This encourages efficient implementation on the part of the computer department, and provides the user with a reliable cost estimate on which to base his implementation decision. One difficulty with this approach is that at the beginning of a project the user's requirements, and the difficulties of implementation, are often not clear. A fixed-price agreement can thus open the computer department to considerable risk, and may well result in considerable waste of time and effort in attempts at renegotiation and arguments as to what was and was not included in the original quote. One way of reducing these dangers is to begin any implementation project with a pilot study whose objective is to develop the design of the proposed system to a sufficient extent that a firm quote for full implementation can be given with confidence (with the user having the option to abandon the project at that point).

Capitalizing System Development Costs

System development costs can be charged directly against users' operating budgets on the same basis as processing charges; alternatively, they can be charged to users' capital accounts and depreciated over the anticipated life of the system. Since system development represents an investment of resources made with the expectation of long-term benefits, capitalization will generally be a more appropriate treatment, unless development is on such a small scale that capitalization is not worthwhile. If system development charges are treated as operating expenses, the impact of these charges on users' operating budgets is likely to distort their decisions on whether, and when, to initiate major systems projects. The capital budgeting process provides a far more appropriate framework for these decisions: not only does it avoid the distortions mentioned above; it also ensures that the formal cost-justification that is normally applied to other forms of capital expenditure, is also applied to systems projects.[1]

---

[1] Note, however, that the normal practice with most capital projects of reducing all costs and benefits to a single measure of financial return, may not be appropriate in the systems area, where a large proportion of benefits are often intangible. (See K. E. Knutsen & R. L. Nolan: "Assessing Computer Costs and Benefits", J. of Systems Management, Feb. 1974, pp 28-34.

Special Pricing Problems

Charging for Data Bases[1]

A common organizational data base represents a shared resource with large implementation and operating costs. The way these costs are treated can significantly affect the probability that such a data base is created, and the effectiveness with which it is used.

Organizations often treat the cost of implementing a common data base as a current operating expense. It is more appropriate, however, to treat these costs as capital expenditures, since they are intended to yield long-term benefits. A manager will be inhibited from undertaking a data base consolidation project if it would produce a large drain on his or her operating budget, particularly in an organization which emphasizes short-term results.

A further problem is that organizations often allocate the implementation cost of a data base among its users.

---

[1] This section is largely based on J. C. Emery & H. L. Morgan: "Management and Economics of Data Base Management Systems", published in D. A. Jardine (ed): Data Base Management Systems, American Elsevier/North Holland, 1974, pp. 185-193.

If, say, a corporate division decides to implement a common data base, then the implementation cost should be borne at that level (i.e., as a divisional overhead). Allocating the costs serves no clear purpose, since it is unlikely to alter the original implementation decision. Moreover, such allocation may actually inhibit use of the data base once it is established -- particularly if the allocation is based on usage.

The ongoing operating costs associated with the data base also have to be dealt with somehow. The costs of retrieving and processing information for specific applications should, of course, be charged to the users involved, and can be handled through the normal mechanism for charging for processing. The costs of data collection and storage (both clerical and computer-related) present a more difficult problem. Where the data are widely, and fairly evenly, used throughout the organization, these costs are best treated as an overhead. Otherwise, actual use of the data base (in terms of number of accesses or volume of data transferred) can be measured -- a sampling method will normally suffice -- and costs allocated proportionately.

In many cases users incur significant expense to provide data that is employed elsewhere in the organization. This raises the possibility of a 'reverse charging'

arrangement through which the data center compensates suppliers of data for the costs of their data collection activities, and passes these costs on to the end-users of the data. Such an arrangement might increase the supplier's motivation to provide timely and accurate data, and would ensure that end-users' charges reflect the full cost of the data they employ. However, to our knowledge no organization is as yet operating on this basis.

Multiple Processing Facilities

Special considerations arise where an organization operates several processing facilities. In this situation management will wish to ensure that work is distributed among the facilities in a way that takes account of their comparative advantages (for example, a powerful machine may be best suited to large-scale computation, whereas a mini-computer may be more efficient for interactive work), and at the same time produces reasonably balanced machine loading.

If management pursues these objectives by controlling centrally the assignment of work, it will be desirable to price services to users so that their charges are independent of the machine to which their work is assigned. The common practice of setting rates independently for the

different facilities so that each recovers its costs will not normally provide this machine independence; instead, the installations' costs and revenues should be pooled. Ideally, a universal set of charging parameters should be chosen, with prices for each installation set so that a given unit of "work" costs the same on any machine (so that, for example, a machine with a processor twice as fast as another would be priced at double the rate per CPU-second.) As an alternative to resource-related pricing, output-related pricing, which is completely machine independent, may be used where appropriate.

If, alternatively, users are left free to choose the installation they will use, the pricing structure needs to be designed to motivate the desired distribution of work. Again, independent pricing of each installation will not necessarily achieve this. For example, a new machine is likely to have lower unit costs than older, less technologically advanced, installations. If each installation is priced to recover its costs, the new one will be cheaper and will become congested while other remain underutilized. Clearly the organization's interests would be better served if, say, the old installations were subsidized to bring their prices down to the level of the newer one.

Competing with Decentralized Resources

A similar problem arises from the existence of alternatives to the computer center, such as local mini-computers and outside services. Again, the computer center's price structure should in general be designed so that if users are left free to choose between the center and the alternatives on the basis of comparative charges, a cost-effective distribution of expenditures will result.

It would appear at first glance that all that is necessary to achieve this objective is to price the computer center's services in a way that accurately reflects their cost. In practice, however, the problem is more complex. Organizations frequently find that users are financially motivated to move away from the central installation when management feels this is against the general interest. The result is either an undesirably large flow of work away from the central installation (often leaving that installation with high fixed costs that it cannot recover), or alternatively the imposition of restrictions on users' freedom to choose their source of supply, which generally leaves users frustrated, and may well result in failure to exploit cases in which mini-computers or outside suppliers have real advantages.

One obvious way of reducing this difficulty is to subsidize the computer center so as to make it more competitive, though this decision obviously must take account of wider considerations than those involved here.[1] Apart from this, a major contribution can be obtained through sophisticated pricing by the computer center. Usually the computer center's rates reflect the costs of the generality and flexibility it provides, even though some users may not need these; this can be a prime factor motivating such users to choose specialized alternatives.[2] Therefore this generality and flexibility should be 'unbundled' (i.e., charged for separately) within the center's pricing scheme as far as possible.

Two techniques in particular are important here. First, major resources such as software libraries and user

---

[1] This possibility is often rejected out of hand by those who feel that a computer facility must recover its costs. In our view there is nothing magic about full cost recovery: the issue is whether a subsidy will produce more or less desirable results.

[2] A typical example is that of the university physics department, which needs large quantities of computational power but no software packages or support services. Such a user may well find it cheaper to set up its own installation, even though the computational resources required can be provided more cheaply by the center.

support services should be charged for directly rather than being bundled in with processing rates: the greater the unbundling of prices, the more each user's charges will reflect only those resources used for his particular application. Second, large-scale users should be offered long-term, minimum-commitment contracts at discounted rates that exclude the extra costs of handling small-scale users and those with variable usage. This allows the center to compete with dedicated mini-computer installations on equal terms.

Chapter 5

ADMINISTRATIVE AND CONTROL PROVISIONS

Implementation of charging involves developing computer accounting routines and associated administrative procedures. We discuss below some of the provisions needed within these routines and procedures to support decision-making and control by those involved in the charge-out system -- i.e.:

- computer center management
- line managers in user areas who have responsibility for computing budgets
- end-users of computer services

Our discussion is oriented toward the job-shop environment typical of academic or engineering and research computing, since this type of environment normally requires more elaborate administrative and control provisions than are needed in an installation limited to routine administrative data processing. However, many of the points equally relevant to the latter type of situation.

Computer Center Management

Since charging involves measurement of resource usage, the charging routines can produce as a by-product information on the overall installation workload and utilization. Such information can be used in fine-tuning the installation's hardware configuration and software, and in balancing workload through pricing and other control mechanisms. Specific information which should be available to computer center management includes:

- utilization reports for the various system resources, showing overall utilization for each resource and breakdowns by class of work (e.g. batch, timesharing, etc.).
- analyses of installation workload by time of day, type of work, priority, and user category.

Breakdowns will be required in terms of revenues as well as resource units, to aid in pricing and revenue forecasting.

A further consideration where charges form the basis for cash transactions is the need for provision to audit the charge-out system.

Budget Administrators

Individuals responsible for computing budgets in the user area need a means of monitoring and controlling the activities within their area of responsibility. This may include not only those with formal budgetary responsibility, but also subordinates such as a project leader, or a faculty member controlling an allocation for instructional use within his or her class.

Budget administrators should be able to set up separate accounts to cover different users and/or activities under their control, and should have available usage reports covering these accounts, including:

- reports showing expenditure for each user or activity, with comparisons against budget, and with the ability to obtain a breakdown to the level of individual jobs if desired.

- reports which show how users are spending their allocations (e.g., by type of processing, by source of supply, and by priority level).

Ideally, these facilities would be provided through a flexible report generator which could be used to aggregate the detailed charges in various ways as needed.

In addition, budget administrators need to be able to control the consumption of users for whom they are responsible. Control provisions might include, for example:

- the ability to limit total expenditure
- where charges vary with priority requested or time of day, the ability to set maximum priority and/or limits submission to off-peak periods.
- a means of limiting resource usage through restrictions on program size and run-time, disc storage quotas, etc.
- a means of limiting type of computing performed (e.g., batch only, or running of 'canned' programs only).
- the ability to restrict the user to a specific source of supply (e.g., main computer center only) when computing funds can generally be used more widely.

Ideally, it should be possible to set these parameters for each account independently. This enables the budget administrator to tailor the restrictions to each user

or group of users. Moreover, where a given user has access to two or more accounts established for different purposes, it may be possible, through restrictions on the accounts, to limit usage on each account to the purposes intended.

End-Users

Individual users obviously also need to be able to keep track of their own expenditures. Thus the system should, for each job or time-sharing session, inform the user of the associated charges. This information should show how the charge is arrived at as well as the total figure, so that the user can explain changes in his charges, experiment with the effects of alternative technical strategies, etc. The system should also provide the user with an indication of the cumulative charges and remaining allocation in his or her account; this information may be provided along with the job charges or on request.

Chapter 6

CONCLUSION

We conclude by returning to our statement at the beginning of this paper: to be fully effective, a charge-out system needs to be tailored to the objectives it is to serve and the circumstances within which it will operate. In our experience, charge-out systems are often less valuable than they might be because management has introduced charging as a 'good idea', without ever clearly defining precisely what it wishes to achieve thereby. Charging is all too often regarded as a straightforward accounting mechanism, rather than a control tool that can can, and should, be tailored to management's needs.

As we have tried to show in this paper, there is a large range of options, both in defining the framework within which charging operates and in designing the pricing scheme itself. Managers should consider whether the particular charge-out system employed within their organization is best suited to its current pattern of computing activities and to their own objectives in relation to controlling these activities. For example:

- If a 'funny-money' allocation system is in use, is it still appropriate? Has it become a focus

for political tensions? Does demand for computing seem insatiable, with no indication as to where the growth in computer expenditure should stop? Is all computing still done at the central facility, without examination of alternative sources such as mini-computers for some activities? If so, a switch to real money charging might greatly ease these problems.

- Is the computer center under-utilized or overloaded? With funny-money budgeting, this may indicate that the total allocation of 'computer dollars' is too small or too large. With real money budgeting, it may be that too rigid an insistence on full cost recovery is preventing the computer center adjusting prices to match demand with capacity.

- Do users ignore charges, fail to look for ways to reduce them? If so, the pricing scheme may be too complex; a simpler, less accurate basis for charging may be preferable. Perhaps the charging system does not provide users with the information needed to enable them to analyse charges effectively.

Or perhaps more effort is needed to educate users in how to interpret charging data. Would transaction-based pricing give better results?

- Are users able to reduce their processing costs by moving their work from the central computer facility, even when this is against the organization's interests? Is this causing dissatisfaction in the case of users who are not permitted to move, and/or loss to the organization in the case of those who are? Would greater unbundling help to bring these users' charges closer to the true costs of their processing? Or more flexibility in pricing (e.g., long-term bulk supply contracts priced at marginal cost for those users large enough to obtain their own machines)? Or should a general policy of subsidizing the center be adopted, abandoning the principle of full cost recovery?

- Are there usage categories whose exclusion from charging would be justified by savings in overhead and red tape (e.g., students, who may represent 50% of total users and 60% of jobs processed, but only 10% of resource utilization)?

- Would competing pressures on the computer center be reduced, and ability to respond to urgent processing requirements improved, through introduction of time-of-day and/or priority-related price differentials?

Readers will no doubt be able to identify similar problem areas associated with computing within their organization. We are confident that a review along the lines indicated above, seeking changes to the charge-out system that might improve its effectiveness, would prove well worthwhile.

FURTHER READING

"Charging for Computer Services," EDP Analyzer, July 1974, Canning Publications, Inc.

Dearden, J. & R. L. Nolan: "How to Control the Computer Resource," Harvard Business Review, Nov.-Dec. 1973, pp. 68-78.

Emery, J. C.: "Problems and Promises of Regional Computer Sharing," in M. Greenberger et al. (eds): Networks for Research and Education, MIT Press, 1974, pp. 189-198.

Emery, J. C. & H. L. Morgan: "Management and Economics of Data Base Management Systems," in J. A. Jardine (ed): Data Base Management Systems, American Elsevier/North-Holland Publishing Co., 1974.

Gill, S. & P. A. Samet: "Charging for Computer Time in Universities," Computer Bulletin, Vol. 13, No. 1, (Jan. 1969), pp. 14-16.

Gladney, H. M., D. L. Johnson & R. L. Stone: "Computer Installation Accounting," IBM Systems Journal, No. 4, 1975, pp. 314-339.

Goldstein, D., M. C. Jensen & D. Smith: Report of the President's Committee on Computing Problems and Opportunities, University of Rochester, 1973.

Grampp, F. T.: "A Computer Center Accounting System," Proc. AFIPS Fall Joint Comp. Conf., 1971, pp. 105-114.

Grillos, John M.: "Pricing EDP Resources," Computer Decisions, (Nov. 1974), pp. 16-17.

Grindley, Kit: "Internal Charging for Computer Services," Accountancy (Great Britain: March 1973), pp. 32-35.

Grobstein, D. L. and R. P. Uhlig: "A Wholesale Retail Concept for Computer Network Management," Proc. AFIPS Fall Joint Computer Conference 1972, pp. 889-898.

Hootman, J. T.: "The Pricing Dilemma," Datamation, August '69, pp. 61-66.

Kantner, H., A. Moore & N. Singer: "The Allocation of Computer Time by University Computer Centers," Journal of Business, Vol. 41, No. 3, (July 1968), pp. 375-384.

Kreitzberg, C. B. & J. H. Webb: "An Approach to Job Pricing in a Multi-programming Environment," Proc. AFIPS Fall Joint Computer Conference, 1972, pp. 115-122.

Lehman, M. M.: "Computer Usage Control," Computer Journal, Vol. 16, No. 2, (May 1973), pp. 106-110.

Li, D. H.: Accounting for Costs of EDP Service Centers, Cost Accounting Standards Board, 441 G. St., Washington, D.C., 20548, 1975.

McFarlan, F. W., R. L. Nolan & D. P. Norton: Information Systems Administration, Holt, Rinehart & Winston, Inc., 1973, Chapter 11.

Management Guidelines for Cost Accounting and Cost Control for ADP Activities and Systems, Report of Recommendations of the Task Group on Principles, Standards and Guidelines for Management Control of ADP Activities and Systems, U. S. General Accounting Office, September 1975.

Marchand, M.: "Priority Pricing with Application to Time-shared Computers," Proc. AFIPS Fall Joint Comp. Conf. 1968.

Marchand, M.: "Priority Pricing," Management Science, Vol. 20, No. 7, (March 1974), pp. 1131-1140.

Nielsen, N. R.: "Flexible Pricing: an Approach to the Allocation of Computer Resources," Proc. AFIPS Fall Joint Conf., 1968, pp. 521-531.

Nielsen, N. R.: "The Allocation of Computer Resources--Is Pricing the Answer?" Communications of the ACM, Vol. 13, No. 8, (Aug. 1970), pp. 467-474.

Nielsen, N. R.: "Using Your Computing Resources to Best Advantage," Journal of Contemporary Business, Spring 1972, pp. 35-50.

Nolan, R. L.: "Preliminary Ideas on Research Design to Investigate Internal Pricing of Computer Resources for Management Control," in Proc. of the Wharton Conference on Research on Computers in Organizations, published as Data Base, Vol. 5, Nos. 2,3 & 4 (Winter 1973).

Nolan, R. L.: "Panel Session: Charge-out Systems for Management Acceptance and Control of the Computer Resource," Proc. AFIPS National Comp. Conf., 1974, pp. 1013-1016.

Nolan, R. L.: Management Accounting and Control of Data Processing, National Association of Accountants, (forthcoming).

Popadic, R. P.: "Design of Chargeout Control Systems for Computer Services," in W. McFarlan & R. L. Nolan (eds): The Information Systems Handbook, Dow Jones-Irwin, 1975.

Reichardt, Karl E.: "Capitalizing Costs of Information Systems," Management Accounting (April 1974), pp. 39-43.

Rettus, R. C. & R. A. Smith: "Accounting Control of Data Processing," IBM Systems Journal, Vol. 11, No. 1, (1972), pp. 74-92.

Schaller, C.: "Survey of Computer Cost Allocation Techniques," Journal of Accountancy, June 1974, pp. 42-46.

Selwyn, L. L.: "Computer Resource Accounting in a Time-shared Environment," Proc. AFIPS Fall Joint Comp. Conf., 1970, pp. 119-130.

Selwyn, L. L.: "Computer Resource Accounting and Pricing," Proc. 2nd. Annual SIGCOSIM Symposium, Oct. 1971, pp. 14-24.

Sharpe, W. F.: The Economics of Computers, Columbia Univ. Press, 1969, Chapter 11.

Singer, N., H. Kanter & A. Moore: "Prices and the Allocation of Computer Time," Proc. AFIPS Fall Joint Comp. Conf., 1968.

Smidt, S.: "Flexible Pricing of Computer Services," Management Science, Vol. 14, No. 10, (June 1968), pp. B-581-600.

Smidt, S.: "The Use of Hard and Soft Money Budgets and Prices to Limit Demand for a Centralized Computer Facility," Proc. AFIPS Fall Joint Comp. Conf., 1968, pp. 499-509.

Stefferud, E.: Untitled paper on wholesale-retail charging schemes, in Proc. SIGMETRICS Technical Meeting on Pricing Computer Services, Nov. 20-21, 1975, published in Performance Evaluation Review, Vol. 5C, No. 1, March 1976, pp 31-70.

"The Effects of Charge-back Policies", EDP Analyzer, Nov. 1973, Canning Publications Inc.

Wiorkowski, G.K. & J.J. Wiorkowski: "A Cost Allocation Model," Datamation, Aug. 1973, pp. 60-65.

DISTRIBUTION LIST

Department of the Navy - Office of Naval Research

Data Base Management Systems Project

Defense Documentation Center (12)
Cameron Station
Alexandria, VA 22314

Office of Naval Research (6)
Arlington, VA 22217

Office of Naval Research
Branch Office, Chicago
536 South Clark Street
Chicago, Illinois 60605

New York Area Office
715 Broadway - 5th Floor
New York, NY 10003

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
(Code RD-1)
Washington, DC 20380

Office of Naval Research
Code 458
Arlington, VA 22217

Mr. E. H. Gleissner
Naval Ship Research and
Development Center
Computation & Mathematics Dept.
Bethesda, MD 20084

Mr. Kim B. Thompson
Technical Director
Information Systems Division
(OP-911G)
Office of Chief of Naval Operations
Washington, DC 20350

Professor Omar Wing
Columbia University
Dept of Electrical Engineering
and Computer Science
New York, NY 10027

Office of Naval Research (2)
Information Systems Program
Code 437
Arlington, VA 22217

Office of Naval Research
Code 102IP Branch Office, Boston
495 Summer Street
Boston, MA 02210

Office of Naval Research
Branch Office, Pasadena
1030 East Green Street
Pasadena, CA 91106

Naval Research Laboratory (6)
Technical Information Division
Code 2627
Washington, DC 20375

Office of Naval Research
Code 455
Arlington, VA 22217

Naval Electronics Laboratory Center
Advanced Software Technology Division
Code 5200
San Diego, CA 92152

Captain Grace M. Hopper
NAICOM/MIS Planning Branch
(OP-916D)
Office of Chief of Naval Operations
Washington, DC 20350

Bureau of Library and
Information Science Research
Rutgers - The State University
189 College Avenue
New Brunswick, NJ 08903
Attn: Dr. Henry Voos