Computer procurement policies for universities and similar users

M. H. J. Webb

London School of Economics and Political Science, Houghton Street, London WC2

Common assumptions about computer procurement are that it is best to get the largest central processor possible, if necessary at the expense of peripherals, that enhancements of existing machines are very cost-effective and that it is better to buy rather than rent. These assumptions appear themselves to be based on assumptions that there are significant economies of scale in computing but relatively insignificant increases in cost-effectiveness over time.

It is suggested that economies of scale may not be realised in computing practice and consequently that a plausible case can be made out for buying balanced computer configurations and rarely enhancing them. If cost-effectiveness is increasing as quickly as some studies have suggested, the financial case for rental rather than purchase may also be very strong. Further research is advocated.*

(Received October 1971)

It has been a common practice, in universities at least, to purchase computers with minimal peripherals and central processors as powerful as possible, and to enhance these machines progressively. This practice can be justified as being the way of obtaining most computing power for a given sum of money if it is assumed that there are substantial economies of scale in computing and that the rate of increase in the power obtained for a given sum of money is relatively low. Such empirical evidence as is available has been surveyed by Sharpe (1969, Chapter 9), and, summarised, it suggests that not only have there been substantial economies of scale, but also that the rate of increase in cost-effectiveness has been very high. Thus, the empirical evidence suggests that one of the two basic assumptions which justify existing procurement policies is invalid, although it supports the other. It is the purpose of this paper to show that both the main assumptions underlying procurement policies, and therefore the procurement policies, may reasonably be questioned.

There appears to be little or no evidence that economies of scale are actually realised in university computing in the UK. Broadly, capital and operating costs appear to move in sympathy with the computing performance available. Some data are given in Appendix 1 which support this assertion. This is surprising, because one of the 'laws' of computing for which there is most evidence is that computing power increases as the square of the cost (sometimes known as Grosch's Law).

Two qualifications are often made to Grosch's Law:

1. At any given level of technology there is a natural maximum to the size of machine which conforms to the law and the performance of larger machines is relatively disappointing.
2. Performance also increases with the progress of technology, resulting in increasing cost-effectiveness with time.

Neither of these qualifications appears to offer a complete explanation for the lack of evidence of realised economies of scale already noted.

Some part of the relatively disappointing performance of large machines must be ascribed to the extent to which each computer has had to cope with all varieties of work. The directors of the computing centres providing a service to London University believe that a large machine, when tuned to a specialised load, will always be more cost-effective than a small machine. The sentiment is clear and credible, even if 'specialised' lacks operational definition. However, there are a number of factors which limit the degree of specialisation of the load on any particular machine.

1. When only one machine is available, it is naturally expected, as far as is possible, to cope with all requirements.
2. Even where several machines are available, different types of task are generated at different times and are associated with different service requirements. Thus, most long computing jobs are believed to be production jobs which do not require rapid turnround. Many short jobs are believed to be for purposes such as program development, for which rapid turnround is believed to be efficient as well as convenient for the user. The tendency is, therefore, to give preferential treatment to short jobs during the day. More short jobs occur during term than vacation, with the long computing jobs illustrating the opposite generation pattern.
3. The same stream of work may be associated with a variety of computing tasks: e.g. small program development jobs often precede long production jobs, and in future it is likely that much development work will, if facilities are available, be carried out on remote personal terminals under multi-access regimes. Further, files may only be available on one machine, and the jobs making use of a single file may be very varied.
4. Users are generally reluctant to change machines unless the advantage to them is clear. This reluctance often appears to be matched not only by nuisance to the user following a change of machines, but also by inefficient machine usage associated with learning how to use another machine.

Thus, there are many reasons for the general purpose use of machines.

There also appear to be a number of other factors which might tend to cloak any inherent differences in cost-effectiveness of different sized machines.

1. Much of the operational work load, e.g. booking jobs, feeding punched cards or paper tape, separating and allocating output to the appropriate user, is not likely to show substantial economies of scale.

---

*The arguments presented in this paper are based primarily upon opinion (hopefully, informed opinion) and only to a small extent upon rigorously established evidence: this reflects a general state of ignorance about the subject. The author is grateful to those who have commented on a draft of this paper, and particularly to the directors of the various London University computer centres for their help and time spent discussing the subject. The author alone is responsible for the contents.

2. It has been the practice, dictated largely by a desire to make short funds stretch as far as possible, to purchase machines with central processors as powerful as possible with minimal configurations of input, output and storage facilities. It has not been possible to reserve these machines for a job-mix for which their limited configuration is best adapted, but they have had to service a general purpose load. It is a plausible hypothesis that these minimal configurations are less adaptable to a variety of tasks than are machines with more input, output and storage facilities.

3. There is considerable evidence that the efficiency of the software associated with particular machines may increase substantially over time. At best this introduces a superficially random element into relative cost-effectiveness. At worst, the latest big machines with minimal configurations and crude software may look unattractive when compared with older, smaller computers with developed software. For example, improvements in software have been responsible for trebling (approximately) the capacity of the University College (London) IBM 360-65 as judged by performance on the same batch of jobs since it was first delivered.

4. There is, in addition, some possibility that, even when developed, software efficiency declines as the machine size increases, and as the complexity and variety of the tasks undertaken increases. In particular, it is sometimes asserted that it is virtually impossible to monitor the performance of some of the larger machines and machine complexes in any detail. An example of the effects of complicating the tasks that the operating software is called upon to manage on a medium sized machine may be quoted from tentative and preliminary observations of the performance of the 1905E at Queen Mary College. Addition of multi-access work to a simple batch stream appears to increase the proportion of the total CP time required by the executive system by 5%, and increase the proportion of total CP time that is not utilised by 10% (i.e. from 10 to 15%, and 5 to 15% respectively).

However, so long as there is a reasonable demand for increases in computing power to enable larger tasks to be tackled, then there is a case for procuring larger machines when they are available; this case is strong if, despite all the difficulties cited, the cost-effectiveness of the larger machine is no lower than that of others.

'Unbundling' of software charges may reinforce the case for large machines, since the charge for each item of software used on each machine may be constant, irrespective of its level of utilisation; the alternative, for those machines to which 'unbundling' applies, may be a high degree of specialisation to limit the range of software required for individual machines.

These arguments all suggest that there might be a large bonus in terms of machine cost-effectiveness if different machines were devoted to specialised job streams. Users would, however, suffer. For example, many small jobs would have to be given slower turnround to maintain steady specialised loads. Forced transfers between machines, which would often be administratively as well as technically inconvenient (e.g. the transfer of magnetic tape files from one machine storage area to another) would occur more frequently; the resulting inefficient usage might be significant.

Furthermore, to facilitate the loading of each machine with specialised work, and exploit the capacity of large machines devoted to specialised work, jobs would have to be collected from a wider area into each machine. There would, notionally, be a choice between a number of separate data transmission, or courier, systems with many institutions being connected independently to several machines, or transmission networks, probably incorporating message switching facilities, linking the various centres. The costs or user inconvenience associated with either type of transmission system would probably be substantial. A network or message switching system might have the effect of complicating the machine operating software systems that part of the purpose of specialisation had been to simplify.

The arguments for a policy of machine specialisation are therefore controversial on cost-effectiveness as well as service grounds. Nevertheless, there may be limited degrees of specialisation which offer advantages, albeit limited, without the disadvantages of complete specialisation.

The very nature of large jobs compels users to accept a lower level of service, particularly in respect of turnround. Specialised use of the largest machines for the largest jobs, as in the proposed regional centres, would probably offer opportunities for tuning the machine configuration to a specialised load, simplifying the operating systems and possibly for limiting the range of software required to be immediately available. There would be little user inconvenience providing files and programs could be transferred easily between the machines used for preparatory or development tasks and the production machine.

Removal of the large jobs to a specialised system would give much greater freedom of choice of the size as well as type of machines required for other purposes, so long as program and file compatibility with the big machine were maintained. (In practice, of course, this condition would often only be fulfilled by smaller machines of the same series.)

It could be that devotion of one or more machines to the specialised function of providing multi-access facilities for elementary teaching and course work usage during terms for, say, 12 hours per day would also be highly attractive. These services will almost certainly have to be limited in scope and availability (like laboratory use by undergraduates) and may be controlled to make only slight demands for easy job transfer between machines. Furthermore, it should be noted that most commercial multi-access services are provided on the basis of dedicated machine (or machine complex) use, presumably reflecting a belief about the relative efficiency of such arrangements.

The cost-effectiveness of dedicating one or more machines to multi-access operation for elementary teaching and course work may well depend upon finding a suitable alternative load to ensure adequate and efficient utilisation when teaching and course work is virtually non-existent—e.g. overnight and in vacations.

Thus, the case for machine specialisation need not stop at the segregation of long production jobs. Equally, however, the arguments for general purpose machines presented earlier are most unlikely to be overthrown completely. In particular, the small user will continue to need a general purpose service through the one terminal facility.

If the need for very large-scale computing is supplied in future by specialised machines then there is no need for the other machines to be based on the largest central processor possible merely to enhance the size of job which it is possible to tackle at all. In these circumstances a re-examination of some of the procurement policies that appear to have been taken for granted hitherto may be of interest.

The practice of procuring the largest central processor possible may be prompted partly by a belief in the availability of economies of scale. If, however, exploitation of the processor capacity is hampered by a shortage of peripheral facilities, then the immediate power availability might be increased if the same sum were expended on a smaller central processor with more peripherals. This case is examined, on the basis of a number of special assumptions, in Appendix 2, where it is suggested that balanced configurations will make more power immediately available than unbalanced configurations under conditions which are usually met in practice.

It might be argued, however, that the unbalanced configurations offer greater opportunities for extremely cost-effective enhancements, and that the overall result would justify the immediate shortfall of power. Such arguments depend to a considerable extent upon an assumption, usually implicit, that the rate of increase of cost-effectiveness with the progress of time, usually ascribed to technological progress, is very low.

However, the empirical evidence suggests that the rate of increase of computing cost-effectiveness has been high, and this suggestion accords with any assessment of the relative power and cost of the machines available 10 or 20 years ago with those available today. Further improvements in cost-effectiveness are generally anticipated. If the power obtainable both immediately and over the life of the machines by enhancing an old machine is compared with that which could be obtained by purchasing a new machine when cost-effectiveness is increasing it can be shown as demonstrated in Appendix 3, that there is a limit on the period after purchase that enhancement is worthwhile. This limit is shorter under conditions which might make unbalanced configurations more immediately attractive. It appears likely that either the case for buying unbalanced configurations or the case for enhancing them is bad, and therefore the common practice of buying unbalanced configurations and enhancing them later is probably bad on one count or the other.

The relationship between the maximum period after purchase that enhancement is worthwhile and the proportion of the total value of the enhanced machine represented by the original version is illustrated in the figure. The curves fall into two categories, one for the instantaneous power purchased by the enhancement, the other for the cumulative power over the remaining life of the enhanced machine; curves are given to correspond to a variety of assumptions about the effects of under-spending on peripherals relative to a balanced configuration. It is when peripherals are assumed to have little effect on the availability of power, i.e. when values of b are low, that the period after purchase that enhancement is worthwhile is shortest. These curves are based on, amongst others, assumptions that Grosch's Law holds and that the rate of increase of cost-effectiveness is that reported.

If, as has been suggested earlier in this document, the increase in available power with increase in expenditure is less than would be expected from Grosch's Law, then the case against purchasing unbalanced configurations is stronger and the limits to the period within which it is worth enhancing machines are shorter than has appeared so far or is shown in the figure.

If computing cost-effectiveness is increasing quickly then the case for renting machinery and exchanging frequently for the latest must be enhanced in comparison to the case for purchasing a machine and retaining it for a long period. This is demonstrated in Appendix 4 where it is shown that according to the rentals and rate of increase in cost-effectiveness reported by Sharpe (1969, pp. 220, 254, 273) it would have been cheaper to rent rather than buy IBM and CDC equipment to obtain a given power if machines were changed relatively frequently.

The cases presented in Appendices 2, 3 and 4 are dependent upon many particular assumptions, and the detailed conclusions should be approached with caution. Indeed, all the arguments presented in this paper are heavily dependent upon the underlying assumptions.

The point that is being made is not so much that these assumptions and the arguments based upon them are 'right', but that they cannot, at present, properly be dismissed out of hand. The justifications for current practices are also founded upon questionable general assumptions. The remarkable feature is the extent to which analytical tools have not been turned onto a subject of considerable concern in a quantitative academic milieu. It is all the more remarkable that the explorations have been so limited because exploration of such matters as operating software efficiency, a subject of interest to many computer scientists, can only be undertaken to a limited extent with such limited information about the performance of machine-software combinations.

Progress can hardly be made without a measure of computational requirements and computer output that offers a greater degree of reproducibility and generality than the 'job' or 'CP second'. The characteristics of the jobs submitted to a single machine appear to change, and those submitted to different machines to vary. The ratio of CP seconds of various jobs submitted to different machines is not constant, and in any case the CP second is becoming a less appropriate measure of computer utilisation as the costs of central processors decrease relative to those of peripheral and storage devices.

The crude measures of throughput used are a reasonable justification for the reluctance to accept both the conclusions of such studies as have been carried out and their implications, including those developed in the appendices to this paper. However, any reluctance to envisage the possibility and potential importance of improvements is not so justified.

Attempts to specify units of resource utilisation, or to calculate usage on the basis of salient parameters such as the number of transactions, and the number and spacing of lines printed, have featured regularly in commercial as well as scientific computing history. But such measures normally incorporate assumptions about the independence of the component activities of a job, even if they proceed in parallel, and the independence of each job's requirements from coincident jobs. The myth of the independence of component activities of a job has never survived the realisation of inadequate computer power to support all necessary peripheral transactions, while the myth of job independence is clearly untenable in a multiprocessing milieu, and was strained before multi-processing became common whenever, for example, the balance of short term requirements varied from time to time from the average. Furthermore, the difference between much commercial and scientific computer usage in these respects is probably much less than is commonly supposed; few commercial installations are really able to keep closely to a schedule of specific tasks because a substantial proportion of the load of most commercial machines is of erratic incidence, consisting not only of program maintenance, modification, development and trial jobs, but also, for example, of re-runs and engineering or operational research calculations.

The automatic monitoring, data collection, storage and processing facilities of a computer system, particularly a multiprocessing computer system, offer almost ideal experimental facilities for:

1. Accumulating an almost unlimited amount of information about specified characteristics of production work for empirical statistical investigations.
2. Monitoring semi-controlled experiments in which collections of specially prepared or selected jobs are injected into the operational work stream.
3. Conducting fully controlled experiments in which the machine is devoted exclusively to a selection of jobs showing predetermined characteristics.

Until this, or similar research clarifies the extent to which there are economies of scale in university computing and the case for or against machine specialisation, it appears that there is room for much greater flexibility about the arrangements for procuring computing facilities for universities. In particular, the general tendency to buy the smallest configuration of the largest possible machine and enhancing it over a prolonged period appears to be very questionable.

Appendix 1

Costs and throughput capacities of university computers

The directors of the computing centres offering a service to London University made the judgements recorded in **Table 1** about the relative powers of different machines. The average number of jobs processed per hour, where recorded in the annual reports of the appropriate centres for the year 1969–70, is also given in Table 1.

Table 1 Relative capacities of different computers

	DIRECTORS				AVERAGE
	a	b	c	d	JOBS/HOUR
Machines					
Atlas	?	1	1	1	—
1905	0·1	0·75	0·5	0·5	22
360/65	1	1	1	1	22
6400	1	1	1	1	—
6600	3	3	3	3	80

The typical cost of a 1905 is around £250,000 (that at QMC cost approximately £400,000), of the UC 360/65 £450,000 (now enhanced to total spend of £750,000), a CDC 6400 £1 million and a 6600 around £1·86 million (but the London University machine cost £2·25 million).

The current expenditure on the QMC 1905 in 1969–70 was £71,000, and on the CDC 6600 (Centre only) was £281,000 (includes only 9 months three-shift operation with a limited service).

The Newcastle/Durham 360/67 (approximately equivalent to the 360–65) cost £810,000 and has a recurrent grant of £126·6K, as against the £360K recurrent grant to the CDC 6600.

It should be noted that capital and operating expenditures are not collected and published on a uniform basis, and consequently these figures are not fully comparable.

Appendix 2

The case for buying a balanced configuration of a smaller machine rather than an unbalanced, cut-down version of a larger machine when the sum available is fixed

Definition

For any given set of tasks and any computer main frame there exists a *balanced configuration* which offers the lowest cost/power ratio for that machine.

Assumptions

1. That the relationship between the cost, c, and power, p, of balanced configurations is given by
$$p = Kc^a$$
where K and a are constants, so long as c does not exceed some critical value.
2. That the relationship between the cost, c_e, and power, p_e, of other configurations utilising the same main frame and requiring less expenditure than the balanced configuration is given by
$$p_e/p = (c_e/c)^b$$
where b is a constant, providing c_e is not less than some limiting sum and the best configuration for the proposed applications obtainable for the sum c_e is purchased.

For a given sum, s, the power of the balanced configuration is p where
$$p = Ks^a$$

For the same sum, the power, p_e, of a cut-down version of a machine whose balanced configuration cost would be c, would be
$$p_e = Kc^a(s/c)^b = Kc^{a-b}s^b$$

The power of the balanced configuration would be the greater if
$$p > p_e$$
or if
$$Ks^a > Kc^{a-b}s^b$$
or if $\quad (s/c)^{a-b} > 1$

as, by definition
$$s/c < 1$$
$$(s/c)^{a-b} > 1$$
if $\quad b > a$

The implications can be explored further by putting figures to a and b.

It is widely believed that $a = 2$, according to Grosch's Law; Sharpe (1969, Chapter 9) quotes an amount of empirical evidence which tends to support Grosch's Law. It is suggested in this note that the practical evidence offered by computing in universities in this country does not obviously support Grosch's Law, but suggests, rather, that the power/cost ratio of different machines acquired at different times is remarkably constant.

The suggestion that the power of cut-down machines follows the form
$$p_e = p(c_e/c)^b$$
has not, so far as is known, been suggested before and so possible values of b have not been explored. It should be noted, as a warning, that this form has been adopted primarily for convenience.

Three cases taken from within London University may give some idea of possible values of b.

1. The University College London IBM 360/65.

Recent enhancements have, approximately, tripled the power (i.e. reduced the time taken to run a batch of jobs used for testing purposes to ⅓), for an expenditure of £300,000. The original configuration cost £450,000, but this was with an educational discount of 40% (now 10%) and before price increases. Consequently these figures for the enhancement and the original cost are not on a comparable basis. It has been suggested by Professor Samet, the Director, that a cost of the original machine which could be comparable with the £300,000 would be £900,000. If this is accepted the value of b would be 3·8.

2. The addition of extended core store to link previously separate CDC 6600 and 6400 computers (estimates given by Neil Spoonley, Director of the University of London Computer Centre).

Taken separately, the power (relative to a CDC 6400) would be 4 and cost £2·86 million.

When linked via ECS the power would be 7 and cost £3·12 million.

For this case, $b = 6{\cdot}4$.

3. Conversion of a CDC 6400 to a CDC 6500 computer.

The conversion would add £400,000 to the basic cost of a CDC 6400 of £1 million, and would double the power.

For this case $b = 2{\cdot}06$.

(This could be claimed as an example of Grosch's Law in action, rather than the effect being discussed here.)

The variations in b revealed by these examples is very large, and give grounds for questioning the assumptions.

Nevertheless, all three cases would tend to support a hypothesis that most power would be obtained from a given sum of money if a balanced configuration were bought in preference to a cut-down version of a larger machine.

Appendix 3

Comparison of the advantages of enhancing a machine already owned and buying a new machine

Assumptions

As for Appendix 2, with, in addition:

1. That the annual increase in power purchased for the same money is w.
2. That scope exists for an enhancement of a machine bought t years ago for sum c_e to be enhanced to a balanced configuration costing c.
3. That the only other alternative is to purchase a new balanced configuration (note, this assumes, following the result of Appendix 2, that it is generally preferable to buy a balanced configuration rather than a cut-down version of a larger machine).

The full power of the original machine as a balanced configuration would have been

$$p = Kc^a$$

the power of the cut-down version was, according to the assumptions,

$$p_e = Kc^{a-b}c_e^{\,b}$$

The additional power of the enhancement would be p_f where

$$p_f = Kc^{a-b}(c^b - c_e^{\,b})$$

The sum available for the purchase of a new machine is $c - c_e$. The power of this new machine would therefore be p_m where

$$p_m = K(c - c_e)^a w^t$$

and $\quad p_m > p_f$

if $\quad K(c - c_e)^a w^t > Kc^{a-b}[c^b - c_e^{\,b}]$

or, taking $c_e/c = g$ (where $1 \geqslant g \geqslant$ minimum limit)

if $\quad w^t > (1 - g^b)/(1 - g)^a$

It is clear that, whatever g, there is a time limit beyond which it would not be worth enhancing a machine to obtain immediately available power.

To find this time limit,

$$t \log w > \log [(1 - g^b)/(1 - g)^a]$$

Sharpe (1969, p. 344, footnote) quotes Knight as suggesting that the average annual increase in the power obtainable for a given cost for scientific computing between 1953 and 1966 was 92·5%.

Values of t taking various values for g and b, assuming $w = 1{\cdot}925$ and $a = 2$ (Grosch's Law) are plotted in the figure.

This comparison has, so far, been limited to the incremental instantaneous power purchased. Two other effects might be significant.

1. The life of the machine might date from the purchase of the main frame, making the effective life of the enhancement shorter than that of a new machine.
2. The enhancement might look less attractive if the alternative were increasing the sum available for purchasing another machine, rather than purchasing a balanced configuration for this sum. Both effects would result in the period within which enhancement was considered being shortened.

E.g.

Let the life of a machine from the purchase of the main frame be L.

Then the cumulative power purchased with a new machine is

$$p_m = Lp_m = LK(c - c_e)^a w^t\,.$$

The cumulative power purchased with the enhancement is

$$p_f = (L - t)\,p_f = (L - t)\,Kc^{a-b}(c^b - c_e^{\,b})$$

In this case

$$p_m > p_f$$

if $\quad LK(c - c_e)^a w^t > (L - t)\,Kc^{a-b}(c^b - c_e^{\,b})$

or if $\quad Lw^t(1 - g)^a > (L - t)(1 - g^b)$

or $\quad w^tL/(L - t) > (1 - g^b)/(1 - g)^a$

since $L/(L - t) > 1$ for all values of t greater than zero, then the limiting values of t will be less in this case.

Values of t, assuming the life of the enhancement is limited to the remainder of the seven year life of the original machine and the same values of w and b as before are shown in **Fig. 1.**

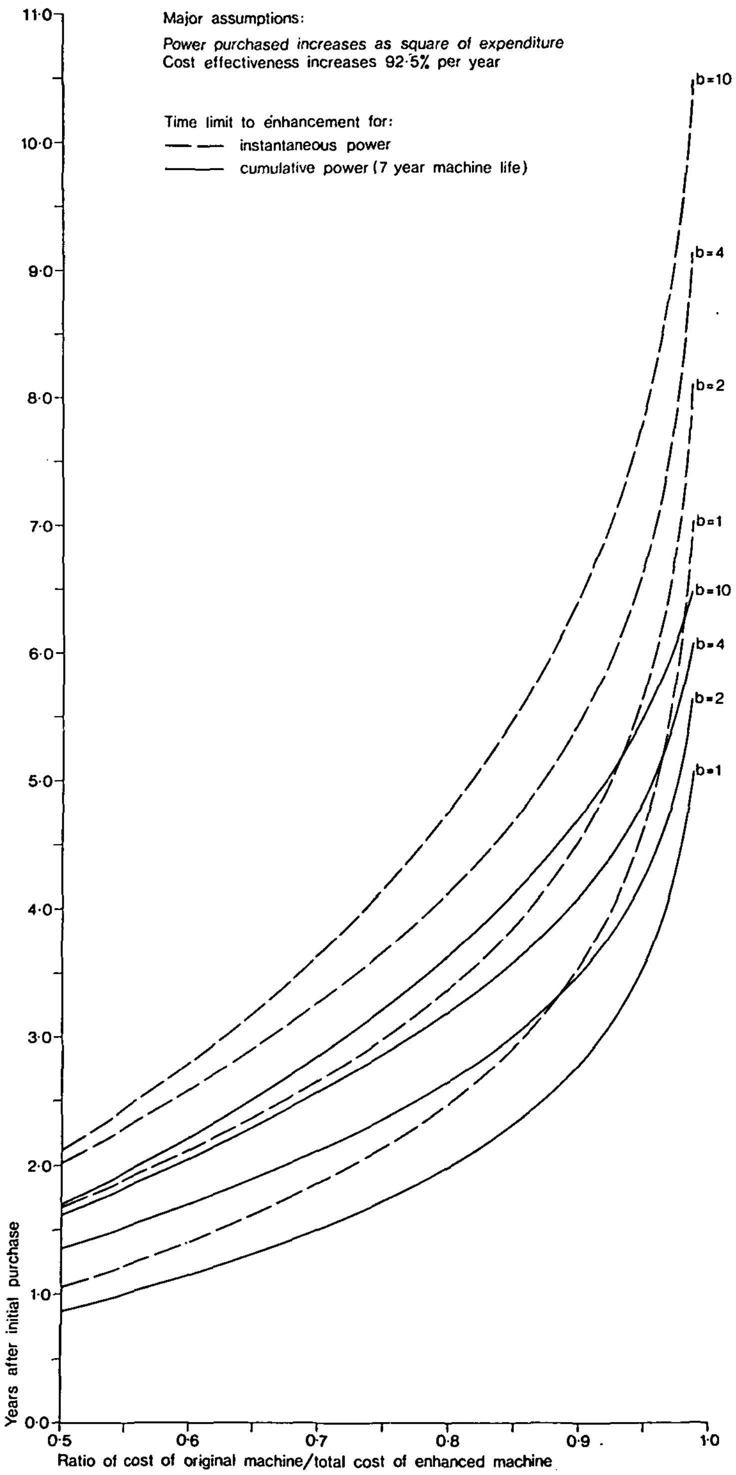

N.B. For derivation see Appendix 3

Fig. 1. **Time limits beyond which enhancement less attractive than purchasing another machine**

Appendix 4

An examination of the case for rental as against purchase

Assumptions

1. Sum c available at commencement of a period to be employed on either rental or purchase.
2. An annual rental charge of a fraction v of the purchase price.
3. An annual decrease in the cost of machines of the same power to a fraction r.
4. An effective machine life of L years, after which it is worthless.
5. A constant annual interest on money deposits of i expressed as a compound interest ratio, *i.e.* $i > 1{\cdot}0$.
6. That the rented machine is changed each year, to take advantage of technological advance (or competition), for a machine of the same power.

For the first year, the rental cost for a machine of capital cost c is vc.

At the end of the first year the sum remaining is $ic - icv$.

For the second year the rental cost of a machine of equivalent power is vcr.

At the end of the second year the sum remaining is

$$i^2c - i^2cv - icvr \ .$$

For the third year the rental cost of a machine of equivalent power is vcr^2.

At the end of the third year the sum remaining is

$$i^3c - i^3cv - i^2cvr - icvr^2 \ .$$

At the end of the $L - 1$th year the sum remaining is

$$i^{L-1}c - i^{L-1}\,cv - i^{L-2}\,cvr \ldots icvr^{L-2}$$

and if

$$r/i = x$$

the sum remaining is

$$i^{L-1}c[1 - v(1 - x^{L-1})/(1 - x)]$$

The rental for the Lth year would be vcr^{L-1} and rental is cheaper than purchase if the sum remaining is greater than this, i.e.

if

$$i^{L-1}c[1 - v(1 - x^{L-1})/(1 - x)] > vcr^{L-1}$$

or

$$1 > vx^{L-1} + v(1 - x^{L-1})/(1 - x)$$

or

$$1 > v(1 - x^L)/(1 - x)$$

or

$$v < (1 - x)/(1 - x^L) \ .$$

The implications can be shown by quoting some general figures for L, v, r, and i.

Sharpe (1969, pp. 344, 353, footnotes) quotes Knight as suggesting that there is an annual decrease in cost of 23·6% for a given performance (scientific computation) and Patrick as suggesting a 23% decrease per year.

i.e. $r \sim 0{\cdot}77$
if i is taken to be 1·10 (10% interest rate, the test rate for government financing)
then x may be approximately 0·7
if L is taken to be 7 and x to be 0·7
then $(1 - x)/(1 - x^L) \sim 0{\cdot}33$.

Turning to Sharpe (1969, pp. 220, 254, 273) again for ratios of rental (excluding maintenance) cost to price for unlimited utilisation in 1966, we have

for CDC equipment

$$v = 12/46{\cdot}36 \times 1{\cdot}2 = 0{\cdot}3106$$

for IBM equipment

$$v = 12/50{\cdot}33 \times 1{\cdot}1 = 0{\cdot}2575$$

Thus by these figures, it would be cheaper to rent than buy both IBM and CDC machines.

Clearly, no account has been taken in these figures of discounts or duty, nor of the inconvenience and cost of frequent changes of machinery.

Reference

SHARPE, W. E. (1969). *The Economies of Computers*, New York and London: Columbia University Press.

Book review

Computers in Number Theory. Proceedings of the Science Research Council Atlas Symposium No. 2 held at Oxford from 18-23 August 1969. A. O. L. Atkin and B. J. Birch (editors), 1971; xvii + 433 pages. (*Academic Press, London and New York*, £8·00)

Many results in the theory of numbers, both before and after Euler's discovery of the law of quadratic reciprocity, have been discovered and confirmed by extensive numerical evidence long before they have been proved mathematically. So it is not surprising that number theorists have turned to modern computers to aid them in their investigations and that they have achieved notable success.

Computers in Number Theory gives detailed accounts of many recent researches in number theory and combinatorial theory. While a few of the papers seem to be on number theory with little or no reference to computers, most are concerned with results in the theory of numbers that have been suggested or discovered or proved by use of computers. Some of the results discovered by computers are given mathematical proofs, others remain uproved, and some of these (for example, those concerning linear relations connecting the imaginary parts of the zeros of the Riemann zeta function) look as though they may remain unproved for a long time. In quite a surprising number of contributions, computers have been used to provide rigorous proofs of difficult mathematical results, where it is by no means obvious, at first sight, that the computer can be of any use at all. A number of these are results in the Geometry of Numbers where conventional methods lead to the consideration of many different cases and subcases all of essentially the same form. It is less surprising that computers are useful in the study of Diophantine equations, but the papers on this topic display great ingenuity in the way the original problems are transformed to forms in which the computer is of use.

While most of the papers are about number theory, there are a few notable exceptions. D. H. Lehmer gives an introductory talk about the economics of number theoretic computation. I. J. Good and R. A. Gaskins develop results in the theory of numbers that are relevant to the art of pseudorandom number generation.

This volume is perhaps best regarded as an essential part of the periodical literature of mathematics. It is of especial interest to number theorists and to those computer scientists who are looking for interesting and unusual ways of using a computer. It has left me with a slight feeling of disappointment, in that, it could, I believe, have been made more interesting and useful if it had included a survey of some of the very striking results in this field that have been obtained over the last 10 years. The present work would have stood up quite well to the comparison with past achievements.

C. A. ROGERS (London)

Errata

Due to delays caused by a strike at the Université de Montréal, the author's proof corrections of the paper 'A note on the generalised Euler transformation' by P. Wynn (this *Journal*, Vol. 14, No. 4, pp. 437-441) were received after the printer's deadline. The following corrections should be made:

On page 438, column two, line 27 '$-\infty < a \leqslant \zeta \pm b < \infty$' should read '$-\infty < a \leqslant \delta \leqslant b < \infty$'.

On page 439, column one, line 11 the equation should be numbered (17).

On page 440, column one, line 12 '*if*' should read '**if**'.

On page 441, line −8 'NORLAND' should read 'NÖRLUND'.

The author's present address is:

Centre de recherches mathématiques, Université de Montréal,
Montréal, Québec, Canada.