Computer Use under a Free-Access Policy

Data on the Dartmouth Time-Sharing System suggest that free access would be inexpensive and practical for others.

Arthur W. Luehrmann and John M. Nevison

There is a long-standing controversy, both within and among universities, regarding the proper method of allocating computer resources to the academic community.

At one extreme, exemplified in varying degrees by most universities, computer use is treated as a marketable good, and allocation is by fee for service, on a pay-as-you-go basis. The ultimate consumer—whether it be a research professor with grant money, an academic department with a budget, or a student with his own funds—checks his pocketbook, then decides how much computing to buy. Open competition in a free market determines both aggregate demand for and individual allocation of computer use.

At the opposite extreme, represented here at Dartmouth College and at a few other universities, computing is regarded as a good that is priceless, in the technical sense in which economists use the word—a good whose subjective worth is extremely difficult for an individual consumer to estimate in advance. Worthless, humdrum hours are punctuated by brief moments of deep satisfaction. In neither state of mind does the perceived worth of computing bear any meaningful relation to its rather easily calculated cost.

Library use is an excellent example of another priceless commodity, and it is so administered at almost every educational and research institution. If the annual cost of the library had to be recovered by means of a borrowing fee, the price would be unbearable—more than $10 per circulation at this university, for example. Such a policy would drive circulation down and increase the borrowing fee still further, until library use finally became the exclusive property of a tiny group of narrowly professional users having grant money or line items in a departmental budget.

Instead, almost all institutions allocate library use by granting free and open access to all members of the academic community. A free-access computer policy is simply an application of the "library model" to computing.

We at Dartmouth have been told for many years by those who favor a fee-for-service policy that our way would lead us to financial ruin, that if we survived it then Dartmouth must be a very atypical university (wealthy, small, nontechnical, not research-oriented, and so forth) and a poor model for others.

In recent months we have been collecting and analyzing data on usage of the Dartmouth computer (*1*), in order to understand better what has actually happened after 10 years of experience with a free-access policy, and to answer some of the questions raised by skeptics. The primary purpose of this article is to bring to the attention of others several surprising displays of actual user data and our analysis of them. In later sections we shall show

1) that the main effect of a free-access policy is that nearly all members of the community use the computer;

2) that, nevertheless, a small fraction of these people account for a very large fraction of the total usage;

3) that this Dartmouth "big user" community is not different from that at other universities and includes students as a small minority;

4) that if one accepts as an inevitable cost the need to supply computer service to these big users, then the added cost of a free-access policy for everyone appears to be no more than a 20 to 40 percent increase in aggregate demand, while the added benefit is a tenfold to twentyfold increase in the size of the total user community compared to the big user community;

5) finally, that a free-access policy does not mean that externally supported projects cannot be charged for computer use, nor that there is no accounting for use.

Before a discussion of the data and their implications, however, a certain amount of background information is needed; first, to define more carefully what our free-access policy is, and second, to show that Dartmouth operates a large computer system with a large aggregate demand similar to that at other universities of comparable size or even larger.

The Unchained Computer

We shall not attempt here to review the history of computing at Dartmouth, which has been told in *Science* (*2*) and elsewhere (*3*). It is enough to say that, from the start, the primary justification for having a computer at all and for increasing its capacity (and the attendant costs) has been the belief that a knowledge of computing would add value to the education of students far in excess of those costs.

Present college students will reach their career and productivity peaks around the year 2000. A recent survey (*4*) indicates that 15 percent of the American labor force today needs to know something about computers. By the turn of the century basic computer literacy is likely to be in even greater demand. In terms of the value of his time, the cost to today's student of acquiring a knowledge of computing is least right now and increases rapidly with each passing year of his life. This sort of "value-added" argument has been persuasive among the faculty and trustees of Dartmouth College. Opinion surveys of recent high school applicants indicate that students also see the promise of learning about computing to be one of their chief reasons for wanting to attend Dartmouth.

Dartmouth College issues to every student, faculty member, and administrator a plastic, wallet-sized identification card with his or her name and an identification number embossed on it.

Dr. Luehrmann is director of Project COMPUTe and Mr. Nevison is administrator of Project CONDUIT at Kiewit Computation Center, Dartmouth College, Hanover, New Hampshire 03755.

That number also represents the individual's personal computer account number, or "user number," in the local jargon. It is all that is needed in order to log into and use the computer (although password protection is also avaliable to each user). Possession of a user number is the right of every member of the academic community.

The final section of this article deals more fully with the crucial topic of budgeting and accounting procedures. Suffice it to say here that an individual's decision to use the computer is not constrained by a concern for costs to be accounted against a personal or departmental dollar budget. By and large, there are no such budget items for computer use at Dartmouth.

This is what free access means to us; and it is essential in what follows to distinguish carefully between free access and unlimited use. Free access means that any individual sitting at a terminal may dial the computer and log in without seeking either funds or permission. It does not mean that he may execute programs that consume arbitrary amounts of computer time or file storage. As the economists tell us, if we forswear money then some other rationing principle must take its place, the total resource being finite.

Rationing of computer resources at Dartmouth is essential and exists in two fundamental forms. The first arises from the nature of a large time-sharing system. On a typical afternoon, when 150 people are in simultaneous contention for machine resources, even the most abusive user would find it difficult to consume more than a few percent of the total resource. Furthermore, as a matter of policy we do not allow any users, with or without grant funds, to buy their way to a higher priority of service, since that would degrade normal service. Thus, time-sharing is intrinsically self-rationing in a way in which batch-processing is not. This is a vital point to which we shall return in our concluding remarks.

The second form of rationing is explicit and requires a modest amount of administrative supervision. A feature of the Dartmouth Time-Sharing System (DTSS), reported more fully elsewhere (*5*), is the ability to establish for each user number a set of specific limits under which that person must work. A typical student, for example, is limited to 32 seconds of central processor unit (CPU) time per job executed, 16,384 words of core-memory during execution, 20,480 bytes of long-term file storage, and no access to the card reader or punch at the computer center. A faculty member has limits also, although somewhat more generous. Note, however, that no user works under a total monthly or annual allotment of computer resources.

Our system would not be workable if the limits were rigid, as there are times when some users need more than their current allotment. The DTSS software makes it a 30-second task to change the limits on an individual or a group, and 10 or 20 such requests are received and evaluated each day by a computer-center staff member, who spends about half her effort in this way. A request for temporarily doubling one's normal execution time limit or file storage limit is usually honored quickly and without deep analysis. Large or permanent changes require formal justification and may take longer or be discouraged.

Our model, once more, is the library. Anyone should be able to browse through a catalog, use the reading rooms, and borrow an armload of books, but not everyone should be admitted to the rare-book collection, and no one should be able to drive a truck up to the loading dock and haul off 10 percent of the entire collection.

Capacity and Aggregate Demand

To become convinced that the usage data presented in the next section are relevant to other universities, one must realize that the size of the Dartmouth computer resource and the total demand for its use are comparable with the figures at other universities. Our description will be brief; readers wishing more detail should investigate the references (*6*).

The computer hardware is a Honeywell 635 dual-processor system with 163,840 36-bit words of core memory and a 1-microsecond cycle time. The time-sharing operating software, DTSS, was written almost entirely by undergraduates with faculty and staff supervision. Program swapping is to two magnetic drums, and file storage is on IBM 2314 disks with a capacity of 400 million bytes. Two communications computers serve as an interface between the central computer and up to 170 simultaneous remote terminals.

Currently, about 400 terminals are available; nearly 300 of them are on the Dartmouth campus in several public teletype rooms, laboratories, faculty and administrative offices, and even dormitories and a few faculty homes. The remaining 100 terminals are used by other institutions in a regional network that includes members in Chicago, New York, Boston, Montreal, and Toronto, as well as a number of small towns scattered across New England.

The operating system, DTSS, currently offers nine programming languages (including Basic, APL, Fortran, and Cobol) and several text editors as part of its interactive time-sharing service. In addition, it offers a "background" service that satisfies the needs usually handled by a batch-processing computer. A user at any terminal may initiate one or more background jobs, typically requiring longer execution times or access to the high-speed printer or magnetic tape handlers; and then he may hang up his telephone and go about his business without further communication with the computer. The majority of administrative data processing goes on in this fashion and can be scheduled to have minimal impact on interactive service.

As we turn now to the aggregate usage figures it should be noted that about 80 percent of system resources are consumed by the Dartmouth users, who number approximately 3400 undergraduates, 600 graduate and professional students, and 400 faculty members.

One standard measure of aggregate usage is the total number of man-hours spent at computer terminals. The terminal time in May 1973 came to 31,499 hours. This amounts to about 1000 terminal-hours per day. For comparison, this is more than three times the May figure for the Multics time-sharing system at Massachusetts Institute of Technology (*7*), with a Honeywell 645 system similar to our 635.

Another standard measure is the total number of CPU-hours actually used, excluding idle time and system overhead. For Dartmouth, the figure for May 1973 was 482 CPU-hours. Because of the similarity in hardware, this number can best be compared with the 375 CPU-hours used by Multics in that month (*7*). Less easily interpreted are the figures of 532 CPU-hours on the Control Data Corporation 3600 and 3800 systems at the University of Massachusetts (*8*), which serve 6200 people, and 162 CPU-hours on the

IBM 360/65 batch system at the University of Chicago (9), with a largely graduate population of about 8000 people.

These comparisons indicate that the Dartmouth computer has a capacity and per capita demand similar to those at other universities of very different characters. Hence, the data and analysis in the next section should not be dismissed because of a presumption of gross differences between systems or levels of demand.

Usage Spectrum

Aggregate usage data tells only part of the story and can be extremely misleading. Our usage data for May 1973 showed 31,499 terminal-hours and about 4000 active user accounts. This might invite one to summarize by saying that the "representative user" uses about 8 terminal-hours per month. As we shall see, such a summary would be a poor way to comprehend the actual data. In fact, the very idea of a representative user is wrong.

Each month the computer produces a summary of all activity for each user account. In any month 4000 to 5000 accounts show some activity. The particular data we have been studying are, for each user account, the terminal time and CPU time used. At first we produced the usual histogram displays showing the number of users in increasing usage bins. After some experimentation, however, we hit on a much more revealing presentation.

Consider the terminal-time data, for example, as a simple list of 4000 or so numbers, one per user. Now, suppose one sorted the list so that small users came first. One might graph the result immediately, but a somewhat better picture emerges if an additional step is taken. We replace each person's actual usage by the cumulative usage due to that person plus all others who had a lower usage than he.

The result for our May 1973 data is shown as the solid line in Fig. 1. Note that the points along the horizontal axis represent all the 4000 users, while the height of the curve above each user represents the fraction of the total usage due to that user and all smaller users, and so goes from 0 to 100 percent. Ascending the curve from the left, one can see at any point what percentage of the total usage was due to what percentage of the users.

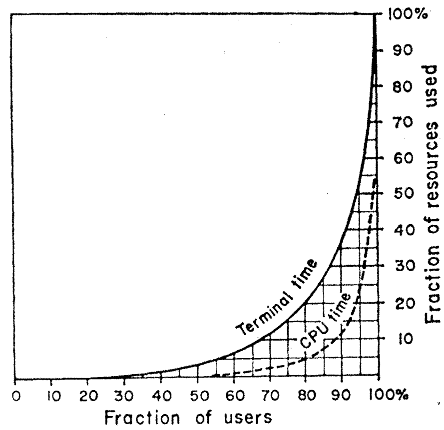

Fig. 1. Dartmouth computer usage data for May 1973. The horizontal axis represents approximately 4000 active users, sorted in order of increasing activity. The vertical axis represents approximately 31,500 hours of aggregate terminal use and 480 hours of aggregate CPU use. At any point on either curve, the height gives the fraction of use due to the fraction of users to the left of the point. Despite their smooth appearance, these are graphs of actual data.

The striking feature of Fig. 1, undoubtedly, is the fact that only a tiny fraction of the total usage was consumed by a large majority of the users. For example, 50 percent of the users (about 2000 people) consumed collectively only 3 percent of the terminal time. Looked at from the other end of the spectrum, the data show that the top 5 percent of the users (about 200 people) used 50 percent of the man-hours spent at terminals.

The dashed curve in Fig. 1 shows a similar graph of cumulative CPU time, for which the skewing is even more evident. Here the 200 big users consumed 75 percent of the CPU time; and the 2000 small users used only 1 percent of the CPU time.

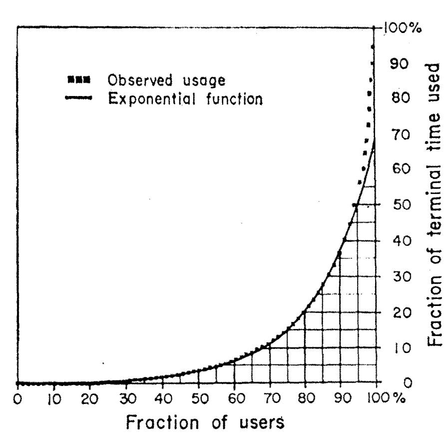

Fig. 2. Least-squares fit of observed terminal-time data to an exponential function. The deviation above 95 percent of the users may be due to the fact that several different people often use the same account number in the case of the most active accounts. Hence, the top 5 percent of user accounts may actually represent more than 5 percent of the users.

It is fair to say that under our free-access policy a substantial majority of our users have a negligible collective impact on the total resource, even though we must assume that each user is satisfying 100 percent of his computing needs. Most users appear not to need very much. Or, to put it another way, obstacles erected to casual computer use would succeed in alienating several thousand people and only gain an additional capacity of a few percent for the "serious" users.

These usage curves fit an exponential function rather well over most of their range. Figure 2 shows a least-squares fit of the terminal-time data to an exponential curve. The consequences of an exponential description of the usage data are noteworthy. Fundamentally, it means that there is a geometrical progression of usage, analogous to progressions of musical pitch or loudness. Only relative values are meaningful, from which it follows that the idea of a representative user is meaningless. To any particular user, a big user will be someone who uses twice as much and a small user one who uses half as much. What we seem to be observing, in effect, is what economists have maintained for decades: everyone's perception of need is in terms of some percentage of what he already has.

Returning to Fig. 1 and noting again that it is the top 5 percent of the big users who consume most of the resource, it is of considerable interest to ask who those 200 people are and what they are doing. Our system of assigning user numbers makes it easy to sort individuals into general categories. For our present purpose we have defined four classes of users: (i) students, including graduate students and students at secondary schools and other colleges, (ii) faculty and research project accounts, (iii) computation center and computer science accounts, and (iv) administrative data processing accounts. Then we have divided the 4000 active users in May into 5 percent intervals, from small users to big users. Each interval contains about 200 users, who have been sorted into these four categories.

Figures 3 and 4 show the results as bar graphs superimposed on the curves

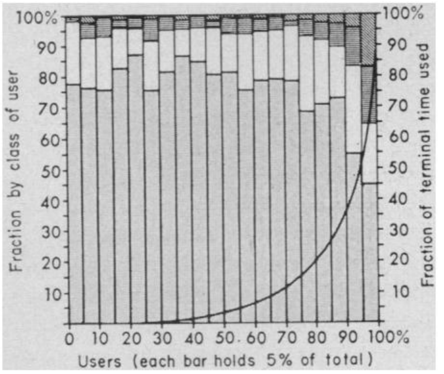

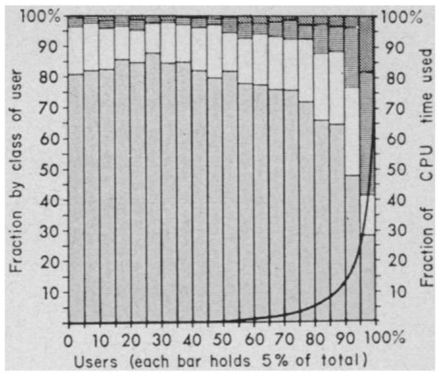

Fig. 3 (left). Distribution of users by class, for terminal-time use. Each 5 percent interval of users (about 200 people) has been sorted into four classes. The bottom shaded area of each bar shows student users. The white area shows faculty and research users. The horizontally ruled area represents computer science and systems development use. The top, diagonally ruled area shows administrative data processing use. The curve reproduces the graph of terminal use in Fig. 1. Fig. 4 (right). Distribution of users by class, for CPU-time use. The shaded bars here have the same meaning as in Fig. 3. The curve is reproduced from the graph of CPU time in Fig. 1. Note that here and in Fig. 3 the small users are mainly students, while the most active users are in the three categories traditionally served by university computers. (Computer science and systems development usage is as large as it is for two reasons. First, most systems work goes on during regular time-sharing hours and so is billed, like any other usage. Second, Dartmouth is the development site of DTSS, a constantly growing and changing operating system.)

shown in Fig. 1. The bottom, shaded area of each bar represents students; the white area faculty users; the horizontally ruled area computer center accounts; and the top, diagonally ruled area administrative use.

Note that about 80 percent of the small-user population are students. Thus, most student use has virtually no deleterious impact on system resources. Looked at in another way, Figs. 3 and 4 support our contention that the big users at Dartmouth appear to be quite similar to the users of a batch-processing, pay-as-you-go computer typical of most other universities. As elsewhere, faculty research, computer science training, software systems development, and administrative data processing dominate the use of our computer. If Dartmouth excluded student use, the cost of operating the computer center and the aggregate demand for service would fall only slightly.

Conclusions and a Caveat

The data given above support the contention that a free-access policy for university computing, far from leading to economic disaster, is a highly practical, inexpensively administered mechanism for satisfying the general computational needs of the majority of students and faculty. The data also support our claim that the Dartmouth computer environment is not unique and that our example might easily be followed by others.

If other universities, currently meeting the needs of the big-user community, should convert to a free-access policy, we venture to predict at least a tenfold increase in the number of users, but only a 20 to 40 percent increase in aggregate demand. Per capita user costs would decrease enormously.

There is a caveat, however. The majority of small users have very small tasks to be performed—ones that require seconds or even tenths of seconds of CPU time. The above conclusions are contingent on having a computer system capable of being used in extremely small quantities by a large number of individuals at approximately the same time. As stated earlier, such a system is inherently self-rationing and prevents the abuses of a free-access policy that would certainly occur if a conventional batch-processing computer were made freely accessible to the university community. (Let anyone familiar with a batch operation imagine the cost of processing 20,000 jobs per day, no matter how small; yet this number is common at Dartmouth.) This is one reason why time-sharing is essential for general university computing, although certain specialized applications may be better served by a batch processor or by minicomputers.

Afterword: Budgets versus Accounts

The idea of unbudgeted computer use and free access is apt to strike terror in the heart of a computer center director or university comptroller concerned with fiscal accountability. They will certainly wonder about the wisdom of leaving the consumer unaware of any measure of his demands for computing resources. They will wonder how a computer center can plan for the future without the budget projections of its customers. And they are sure to wonder how computer use can be charged to externally supported projects while it appears to be given away freely to students and faculty.

In settling these doubts, it is essential to distinguish between budgeting and accounting. Budgeting is anticipatory, while accounting occurs after the fact. It is true at Dartmouth that academic departments do not have budgets for computer time; nor does any course have a computing budget to be dispensed to the students in the course; nor does any dean or committee have a budgeted sum that students or faculty members can appeal to for computer use. Virtually the entire budgeting mechanism for computing is avoided.

Accounting is another matter. All computer use is accounted for, down to the last penny on every one of the 12,000 user numbers active during a year. The DTSS billing program accumulates for each user itemized data on terminal time, CPU time, file storage, lines printed, cards read, cards punched, and so forth. Each month it generates for every user an actual bill showing the dollar cost of each item used. Rates are based on equipment rental and depreciation costs, staff salaries, and building costs. Rates are the same for all Dartmouth users, whether individual students or federally supported research projects. (Off-campus rates are slightly higher.)

With the single exception of students, the monthly bills are mailed to all users. Off-campus users make cash payments. At the bottom of each on-campus bill a printed message states a college account number to which the total amount will be charged unless the recipient takes special action. If the user number belongs to a supported research project, the project will be charged. If it belongs to an individual faculty member, a general unsupported faculty research or education account, supervised by the dean of faculty, will be charged. While student bills are not mailed to the individuals, they too represent actual charges against a general account supervised by the dean for student affairs.

It must be stressed that all user accounts are paid in real dollars. Supported accounts are paid with dollars granted to the college by outside agencies for the support of research. (Computer use is a direct expense in grant budgets.) Unsupported accounts are paid with dollars that come from tuition payments, endowment income, and so on. Since the billing rate, which has been audited by a federal agency, is identical for both classes of account, our free-access policy is consistent with federal regulations as expressed, for example, in *Circular No. A-21* of the Office of Management and Budget. The basic difference between the two types of account lies in the fact that unsupported accounts are not restricted by a predetermined budgetary ceiling. In no sense do students "use up" a budget for computing. The advantages of this accounting scheme over the more usual budgetary approach are many.

1) Students, whose net usage has been shown above to be a small fraction of the whole, are less inhibited and more independent in deciding whether, when, and how much to use the computer.

2) Teachers do not wait for "next year's budget" if they decide to assign computer use in a course or to begin a research project.

3) There are none of the usual budgetary incentives to consume one's entire allotment, for fear of a cutback the next year, or to stimulate a black market or barter economy in computer time.

4) Committees and deans do not waste time attempting to decide the proper amount of computer time allotted to each department.

5) Without dollars budgeted for computing, there is no basis for a department to attempt to convert "computer dollars" into "real dollars" to be spent for other needs.

6) Perhaps most important, faculty members with grant funds have not been able to make the computer center their captive and to distort university computing priorities for their own ends.

These advantages have come without sacrificing the need to account for usage, to bring abuse to light, and to charge externally supported projects for computer use. Approximately one-third of the total budget of the computation center derives from cash income from off-campus users and supported projects.

It can be countered that removing computing budgets from user control may restrict user choice and, in effect, make the user the captive of the computer center. In response, we first point out that in the overwhelming majority of universities, computing budgets now have to be spent at the computer center. These computer dollars are a kind of scrip, redeemable only at the company store. Secondly, and more fundamentally, we argue that there are other powerful and less potentially harmful mechanisms by which users can shape the policies and types of service at their computer center. For example, most universities have a library committee that acts as a user watchdog over library policies and service. Dartmouth has a parallel computer committee for the same reason. Finally, we reiterate that the process of establishing fair and rational budget amounts for computing is itself very expensive and fraught with difficulty and mistrust.

In light of this analysis of our experience in 10 years of free-access time-sharing, a university policy-maker elsewhere might well reconsider the available mechanisms for controlling computer use and give serious thought to providing the entire academic community with computer service modeled after the library service it now enjoys.

References

1. J. M. Nevison, *Who Eats the Oysters?: Representative Usage on the Dartmouth Time-Sharing System* (Kiewit Computer Center, Dartmouth College, Hanover, N.H., 1972).
2. J. G. Kemeny and T. E. Kurtz, *Science* **162**, 223 (1968).
3. R. F. Hargraves and A. Stephenson, *Proc. Am. Fed. Inf. Process. Soc. Spring Joint Comput. Conf.* **34**, 657 (1969); T. E. Kurtz, *ibid.*, p. 649; E. D. Meyers, *ibid.*, p. 673; A. K. Morton, *Comput. Autom.* **19** (No. 3), 2 (1970).
4. *A National Survey of the Public's Attitudes Toward Computers* (American Federation of Information Processing Societies, Montvale, N.J., and Time, Inc., New York, November 1971).
5. J. S. McGeachie, *Proc. Am. Fed. Inf. Process. Soc. Spring Joint Comput. Conf.* **34**, 665 (1969).
6. *DTSS—the Dartmouth Time-Sharing System* (Kiewit Computer Center, Dartmouth College, Hanover, N.H., 1973); *Biennial Report* (Kiewit Computer Center, Dartmouth College, Hanover, N.H., 1974).
7. *Bulletin No. 115* (Information Processing Center, Massachusetts Institute of Technology, Cambridge, June 1973).
8. C. Wogrin, private communication.
9. F. Harris, private communication.

NEWS AND COMMENT

European Community: Pragmatic Is the Word for the New "Europeans"

A year ago, the member countries of the European Community were mildly at odds over such things as their common agricultural policy and financing of regional development programs, but they were mainly occupied with adjusting to the arrival on the scene of Britain, Denmark, and Ireland as the Six became the Nine. Then chronic monetary problems worsened and last winter's energy crisis detonated to shake the Community to its foundations. Now, pessimists see the spirit of cooperation which nurtured the Community seriously eroded.

In recent months, the British Labour party has taken office with a pledge to renegotiate membership in the Community, the Italian government has applied draconian import restrictions in apparent contravention of the Treaty of Rome, on which the Common Market is based, and West German chancellor Willy Brandt has resigned as a result of a security scandal. These events and the death of French president Georges Pompidou have caused a number of observers to note that the present malaise of the Community is really a symptom of the political weakness of the member governments. By recent count, seven of the Nine had coalition governments and the other two, minority governments. Coalition politics is a way of life on the continent, but it is not a formula for maintaining momentum in either national or Community affairs.

Since the new year, the three dominant members of the Community, Britain, France, and West Germany, have seen new leaders take over. The defeat at the polls of Conservative prime minister Edward Heath in Britain, the resignation of Brandt, and the death