

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Annals of Economic and Social Measurement, Volume 3, number 3

Volume Author/Editor: Sanford V. Berg, editor

Volume Publisher: NBER

Volume URL: <http://www.nber.org/books/aesm74-3>

Publication Date: July 1974

Chapter Title: Toward Computer Networking: The Harvard Experience

Chapter Author: Joe B. Wyatt

Chapter URL: <http://www.nber.org/chapters/c10177>

Chapter pages in book: (p. 541 - 550)

COMPUTER CENTER NOTES

TOWARD COMPUTER NETWORKING—THE HARVARD EXPERIENCE

BY JOE B. WYATT*

Economic pressure is a powerful motivational force for change. At least one of Harvard's numerous innovations—namely, the divestiture of a major part of its in-house computing service resources—has been stimulated by fiscal deficit. This particular innovation has worked well for Harvard during its early implementation for several reasons, not the least of which is the stemming of a fiscal hemorrhage which accumulated a \$1.6 million deficit in five years. It is clear that the changes in philosophy and procedure resulting from the solution to this fiscal problem have implications more far ranging than just economics. This assertion is based on three hypotheses which have been supported to some extent by Harvard's experience and for which supportive evidence continues to accumulate.

1. It is not damaging to the basic functions of an educational institution to perform a major part of its computing service using remotely located resources.
2. The point has been passed (if it ever existed) in the development of information technology when a single computer system can best satisfy all of the information processing needs of an educational institution.
3. The potential viability of digital communications networks, coupled with the influence of semiconductor technology and related storage technologies, will substantially alter the economic and behavioral models of information processing.

This paper examines one issue which cuts across the assessment of these hypotheses—the economic incentive of a multi-resource network which could exist for a community of computer users. The sample is small and microscopic. The results, however, are very suggestive if incomplete.

THE SETTING FOR CHANGE

In 1971, after accumulating a deficit of \$1.6 million in seven years, the Harvard Computing Center divested its in-house computer hardware. The XDS Sigma 7 and the IBM 360 model 65 configurations, both rented, were returned to their manufacturers and the staff of the Harvard Computing Center was reduced by 40 percent. Replacement services were arranged from remote computing services separately for time-sharing and batch usage. The functions of the Harvard Computing Center were affected substantially. The user interface function, for consulting with computer users relative to applications and problem solving, remained essentially the same. However, since remote computer services replaced the

* This work sponsored in part by NSF Grant GJ40586.

on-campus hardware, the operations staff and the systems programming staff were virtually eliminated. Remote computing services were used as replacements including:

1. A joint computing service center for batch computing was established with MIT using MIT operations staff and an IBM 370/155 located at MIT.
2. A contract was consummated with First Data Corporation for the use of a DEC System 10 for time-sharing services.

With some modifications, these agreements have remained in place since the Fall of 1971.¹ Since November of 1971, the average usage has been 12,211 jobs per month (standard deviation of 1,778) at an average total cost of \$98,531 per month (standard deviation of \$12,614).

Usage of the joint center by Harvard began in October 1971. By June, the end of the academic year, the joint operation faced an operational dilemma. MIT had installed the IBM Time-Sharing Option (TSO) to the operating system of the 370/155 and was actively encouraging its use. Harvard was relatively satisfied with the First Data arrangement. MIT had utilized approximately 65 percent of the capacity of the joint center's 370/155 and was pressing for an upgrade to a 370/165 configuration. Harvard experienced no similar need for processor capacity. The objectives and the philosophy of the two institutions relative to computing services were sufficiently different that the arrangement could not continue to function as a partnership.

As the solution to this problem, it was agreed that MIT would implement its plan to upgrade to a 370/165 configuration and would continue to operate the Center. Harvard became a major customer of MIT, buying approximately 25 percent of the capacity of the system for a guaranteed minimum price. The IBM 370/165 processor was installed at MIT in September of 1972.

Few Harvard computer users will argue that the incremental effect on the user of replacing the on-campus resources of the XDS Sigma 7 and the IBM 360/165 with access to remote facilities was not noticeable. The program conversion effort from the Sigma 7 to the DEC System 10 was agonizingly complex. Conversion of programs from the IBM 360/65 to the IBM 370/155 was not without difficulty. It is still argued by some that worthy programs were not converted because of the short-term lack of funds. Others argue that the conversion effort had a purifying effect by eliminating programs of lesser significance. In the cursory post facto examination which was made, one fact loomed distinguishable. Except for the relatively small number of time-sharing applications which resulted in significant conversion effort, most of the major conversion problems related to the IBM 7090 to IBM 360 move in the late 1960's, a non-trivial phenomenon but unrelated to the divestiture. Nevertheless, the effect of all of these changes in the computer user community had its effect on user attitudes.

The change from the IBM 7000 series of computer systems to the completely incompatible IBM 360 series produced a national trauma.² Users were urged to spend substantial sums in program conversion in order to achieve the cost/performance "savings" offered by the IBM 360. Most succumbed. Many were adversely affected by the costs in time and dollars for conversion, not to mention the major fiasco created by IBM because of the delivery schedule delay and specification defaults of the IBM 360 Operating System (OS/360).³

At Harvard, change and conversion (from the IBM 7090 to the IBM 360/50 and the XDS 940 to the IBM 360/65 and the XDS Sigma 7) had been justified by the management of the Computing Center on the basis of economic factors: an economy of scale producing greater cost effectiveness to the user. The underutilized capacity and the corresponding price increases to compensate without reduction of expense produced the deficit which triggered the user demands for relief. This demand resulted in the organizational changes, the philosophical changes, and the removal of the major on-campus computing resources. A survey of computer usage at Harvard completed in April of 1973 has revealed that approximately 30 different computer systems are currently being used by the Harvard community.⁴ Approximately half of these systems are not at Harvard. Ten different manufacturers' systems are represented. Twenty-two of the systems are not program compatible at the "machine language" level. Some of the usage is very limited, some is extensive. The "locus shifts," however, are generally directed by improvement in price and performance, not by administrative fiat or by academic discipline groupings.

Harvard now represents diversity in computer usage. There is a balance of on-site and off-site computing resource usage with multiple computer systems involved. Measured in terms of equivalent dollar value, Harvard continues to use off-campus computing facilities to satisfy a majority of its computing needs. Each of these systems is utilized on the basis of an evaluation which included economic comparison of alternatives.

THE MICRO-MARKET

One of the more important factors apparent from the Harvard experience is that with all funds for computer service placed in the hands of the users, there appears to exist a viable market. It is true that university computer users generally have a fixed budget for computing services. However, when true computing resource alternatives become available, some shopping takes place. A simple market is created. On the basis of the limited evidence available at Harvard, the market is sensitive to price and service time (response or turnaround) within a menu of "equivalent" services.

Since the beginning of the joint Harvard-MIT Center, each institution had established its own price structure, creating a "Harvard rate" and an "MIT rate," each based on a component usage profile and each designed to recover total cost to the institution. The MIT rate favored relatively processor intensive usage; the Harvard rate favored relatively input-output intensive usage. The affect of the difference between the Harvard prices and the MIT prices on Harvard users was subsequently analyzed over a period of several months (Table 1). It was thus determined that approximately 10 percent of the jobs processed at the Harvard prices could benefit by reduced cost using the MIT prices. The affect on revenue was estimated to be less than 5 percent. As a result, beginning in March of 1973, Harvard users were allowed to choose between two alternative price structures on an individual job basis. Approximately 1,350 jobs per month representing a total average cost of \$12,000 per month have been processed using the "MIT-like" alternate price structure.

TABLE I
COMPARATIVE EFFECT OF HARVARD PRICES AND MIT PRICES

	Oct. 1972	Nov. 1972	Total
Actual billings at Harvard prices :	166.630	93.900	200.530
Lower values of jobs were :			
Cost at Harvard prices < cost at MIT prices	53.530	44.240	97.770
Cost at Harvard prices > cost at MIT prices	44.990	41.730	86.720
Totals	98.520	85.970	184.490
Variance from actual billings: \$	8,110	7,930	16,040
" "	7.6	8.4	8.0

The Harvard chemists discovered during 1972 that certain of their processor intensive computations could be performed substantially less expensively using the Columbia University 360 model 91 than by using the IBM 370 model 165 at MIT, at an acceptable sacrifice in turnaround time. Some usage of the 360/91 was begun and continues on this basis.

A more general examination of the Columbia alternative was begun in May of 1973. A population of jobs were compared between MIT's 370/165 and the 360/91 at Columbia. The comparison, involving 13,758 jobs, was conducted by taking the accounting summary data from all the jobs actually executed on MIT's 370/165 at Harvard prices for March of 1973 and processing these data against the Columbia price structure. Service priorities calling for differences in rate structure were sufficiently similar that it was possible to make a one-to-one mapping of the five priority categories at Harvard-MIT onto the five priority categories at Columbia. For purposes of the experiment, no extraordinary costs were included; only the basic prices of each institution were used (excluding any communications costs, etc., which would be required to communicate with Harvard users). In the aggregate, the total March workload cost \$109,302 at the Harvard-MIT prices and \$119,529 at the Columbia prices.

However, Figure 1 shows a different perspective of the job cost comparison. In Figure 1, the number of jobs for which the Harvard-MIT cost was at least \$10, a total of 2,538, and for which the total cost for job processing at Columbia rates was less is plotted against the percentage dollar savings for job processing. The number of jobs which satisfy the criteria in Figure 1 is 951, 37 percent of the total population. Slightly over 400 of these jobs would have been processed at a saving in excess of 20 percent using the Columbia rates.

The conclusion reached from this first experiment was that the savings in cost for some individual jobs justified further investigation of the cost and method of providing alternate remote services.

Encouraged by the potential alternative at Columbia, a similar analysis for IBM 360 model 91 at Princeton was performed utilizing the same 13,758 jobs from the March, 1973, usage of MIT's 370 model 165. Since Princeton offers a rate structure not geared to priority service categories, all service categories of usage from the MIT computer system were compared to Princeton's "standard" service rate, presenting some bias for a comparison of price and service response time. As




Figure 1 "Ten Dollar" Jobs With Savings At Columbia Prices

in the case of the Columbia system, cost comparisons used basic Princeton prices excluding communications and related costs which would be experienced by Harvard users. On this basis, the total March workload which cost \$109,302 at the Harvard-MIT prices, cost \$51,092 at the Princeton prices!

A comparison of job costs by Harvard service classes between the Harvard-MIT, Columbia, and Princeton price structures is shown in Table 2. The results

TABLE 2
HARVARD-COLUMBIA-PRINCETON COMPARISON BY HARVARD PRIORITY CLASS

Priority	No. Jobs	Harvard (\$)	Columbia (\$)	Princeton (\$)*
Restricted	10	559	598	1,121
Deferred	3,282	25,528	20,728	13,132
Low	631	5,742	5,959	2,809
Student	8,146	54,696	62,002	23,931
High	435	5,314	5,558	1,243
Advanced	1,253	17,447	24,670	8,853
Emergency	1	16	12	3
Total	13,758	109,302	119,529	51,092

* All at Princeton's "standard" priority and price.

of this comparison indicate that for most priority service categories, the Princeton job costs would be substantially less than Harvard-MIT, ranging from 49 percent to 81 percent by category. Again, a different perspective can be found in those jobs which cost at least \$10 for processing at Harvard-MIT prices and for which costs




Figure 2 "Ten Dollar" Jobs With Savings At Princeton Rates

were less at Princeton prices. A summary of this comparison is shown in Figure 2. In this case, the total number of jobs which satisfied the criteria were 2,529 of the total 2,538! Equally startling in this comparison is the fact that 2,521 of these jobs would experience savings in excess of 20 percent in cost, and 2,280 would experience savings in excess of 40 percent in cost. The results of the Princeton comparison indicate a considerably larger margin for ultimate savings to the user on the basis of the assumption used, which included the bias of applying Princeton's "standard" rate to all jobs regardless of priority (and price) category for Harvard-MIT.

To provide a more microscopic summary to the economic comparisons of the use of these three computer systems and to test some of the assumptions used in the comparisons, a selected set of benchmark problems were processed on all three computer systems. These benchmark problems were selected to represent various types of usage from the Harvard community, and each was chosen from

the set of 951 jobs which were processed at less cost at Columbia prices than at Harvard-MIT prices. Table 3 contains a comparison of the actual cost for processing each of these jobs as computed from the individual price structures of the three systems. It is quite clear from this limited comparison that substantial savings could accrue to users from one type of job to another by "shopping" between the three computing resources. These factors would serve to support the hypothesis that establishing a communications network between Harvard, MIT, Columbia, and Princeton, could conceivably create a rewarding market in computer services.

TABLE 3
ACTUAL COST COMPARISON OF SELECTED BATCH JOBS

Job	Harvard	Princeton	Columbia
06A	\$6.24	\$3.06	\$5.53
07	3.34	4.31	2.73
10A	4.87	1.30	3.61
11	1.81	1.27	0.50
10	24.22	22.94	10.35
13A	34.61	10.44	21.65
17	2.33	2.16	1.11
18	2.35	2.12	1.16
CHM	80.60	46.89	27.79
07A	8.49	4.62	5.49
ASM	6.77	3.50	5.57

Note. All jobs were processed using Princeton "standard" prices, the Columbia "standby priority" prices, and the Harvard-MIT "deferred priority" prices.

THE OBSTACLES TO GENERALIZATION

In the opening paragraph of the paper, it was asserted that Harvard's plan for substantive use of remote computing services might be generalizable. The ARPANET has already proven that such a system is at least in part technologically viable. The issues raised, however, are much greater in scope than a technological discussion. The issues also include questions of organization, economic implications, and political/regulatory decisions.

The role of legal and regulatory developments cannot be underestimated when considering an alternative "scenario" for computing without a local computing center. The Federal Communications Commission has opened the inter-city digital communications market to competition, allowing companies other than AT&T to provide such services. The telephone company has responded with a new service offered called "Digital Data Services" (DDS), based on an economy in multiplexing called "Data Under Voice" (DUV), approved by the FCC for operation in 1974. The rate structure for DDS will substantially decrease some digital communications costs.⁵ Packet Communications, Inc. has been awarded the first license for the commercial operation of a computer communications network to operate nationally using the packet switching concept from the ARPANET. Numerous other companies, including Datran and Microwave Communications, Inc., are already in the digital communications business. Some

of the major market research studies of the commercial remote computing industry have resulted in startlingly optimistic projections. These developments would indicate the availability of a variety of digital communications facilities for establishing multi-resource computing and computer-based information networks.

In order to obtain a feel for the communications problem for such a network, several limited experiments have been performed at Harvard. First, a small number of jobs from the Harvard workload were transmitted to the IBM 360 model 91 computer system at the University of California at Los Angeles using the ARPANET. Batch job input and output were successfully performed using the DEC System 10 in the Center for Research in Computing Technology at Harvard which is connected to the ARPANET. However, it was demonstrated that interfacing to the ARPA network for such operation is a non-trivial endeavor requiring several man weeks in programming. Moreover, the experiment demonstrated that the translation required from packet nomenclature to the nomenclature of the OS/360 spooling queue on input and the reverse on output required sufficient processing by the host computer at UCLA to measurably increase the cost of each job. Considering the current cost of an ARPA IMP of approximately \$100,000 and adding the cost for telephone line communications, the viability of the ARPANET for general usage, although technologically sound, is not obvious economically at "full" cost. It appears from the cost projections for a more contemporary packet switched network that the communications costs would not be prohibitive based on the proposed rate structure of PCI.⁶

Second, an analysis was performed assuming the use of conventional "leased line circuits" from AT&T. It was determined from this analysis that, for this set of experiments, a conventional telephone line represented an economically viable solution operating at 4,800 baud. Typical costs for this communications alternative are shown in Table 4. Parenthetically, by the use of a software package available for the IBM 370 model 145, the spooling to the Columbia computer system could be handled simultaneously both with the spooling to MIT and with the operation of stand-alone DOS programs with no change in the configuration of hardware.

TABLE 4
COST COMPARISON FOR COMMUNICATIONS SERVICE
(MONTHLY) FROM HARVARD

	4.8 KB Circuit	50 KB Circuit
MIT	\$270	\$500
Princeton	910	5,235
Columbia	684	3,670
UCLA	3,055	25,252

Of course, the basic "raw" computing cost including the communication cost does not nearly represent the total cost to the remote user. As a matter of fact, in the two years of the Harvard-MIT arrangement, the sum of these two costs represents approximately *half* of the total costs. The remainder of the cost of remote usage is represented by the operation and management of user terminal facilities and services. If this arrangement is typical, as it is felt to be, any resource network must provide that the "buyer" institution would receive a discount from

the "seller" institution sufficient to cover all of the costs of communications, operation, management, and user consultation. The Harvard-MIT agreement for computer services for 1973-74 is at least a comparative model for such a discount arrangement. MIT sells Harvard a guaranteed minimum amount per year of the "raw" services of the IBM 370/165. Harvard pays all other costs including those for all communications, RJE terminals, interactive terminals, supplies, all user consultation and support, accounting and billing, and installation management for remote usage. MIT grants Harvard a discount of approximately 60 percent computed from MIT's "over-the-counter" local rates. A similar arrangement is being finalized between Harvard and Princeton.

Major benefits accrue to each institution in such an arrangement. The seller receives substantial additional income at little or no increase in expense and at no additional risk. The buyer receives a cost-effective service. The computer users of both institutions benefit from the economy of scale of a large computer system which is justified on the basis of the combined workloads.

These data, of course, represent only an illustrative analysis from one perspective. The data do show rather conclusively that a potentially viable market exists, assuming that computer users hold fully fungible dollars for obtaining computer services, as is the case at Harvard; and that the costs of remote usage, including communications costs, can be more than offset by the savings represented from the availability of multiple resources.

CONCLUSIONS

There has been substantial testimony here and elsewhere to the fact that the basic ingredients exist to implement a more general computer network resource.⁸ It will become necessary, however, that a general technological methodology be developed which would allow users to effectively utilize multiple computer resources. This situation is not yet the case. For example, it is now a non-trivial task to modify the operating system control statements necessary to move a job between similarly configured IBM 370 systems, although each of the systems may be similar in configuration and may operate under OS/360. It is also an unpleasant chore for a user to recognize the differences in conventions and technical minutiae that exists between such resources. Extending this task to the more serious incompatibilities that exist between different manufacturers' computer systems complicates the problem dramatically. Charles Holt⁷ has called for "articulated programming" creating a set of standards for programming languages. However, standardization in this area is not occurring quickly enough to alleviate this problem in a reasonable time, if ever.

In order for a simple multi-resource to be user effective, a methodology to address several significant technological questions must be developed. First, techniques must be developed for "translation" of one operating system control language to another. Translation techniques between differently configured and architecturally different computer systems will be needed to move programs or data prepared in one of the standard languages from one computer system to another without a major effort. Second, techniques must be developed for communicating information from network computer resources to users about the

variety and usage of the services of the resource. Many computing centers are already weak in the area of user consultation or marketing their services. The implementation of remote access networks will focus attention on the necessity for adequate documentation and communication mechanisms between computer users and computer resources. Third, communication systems are required which will allow multiple types of user terminals to be transparently attached to the network facilities. Fourth, a viable method for pricing and accounting for resource and communications facilities must be developed. If this is not imaginatively and efficiently done, the problems created could make the network a dubious opportunity for both buyers and sellers.

These technological issues, not unlike the other issues, range from the simple to the complex. Within the range, it is important to determine where reasonable combinations of solutions exist through the use of technology which will deal adequately with these problems. If reliable and generalizable technical solutions can be implemented and used at modest incremental costs, a major evolution in the usage of network computer resources in colleges and universities will be enabled. Research is needed to understand and develop the alternative technological and economic methodologies presented here. Experiments will be required involving multiple resources and multiple users to understand the operation and use of general purpose networks. If such a set of methodologies can be implemented as hypothesized, a new set of opportunities will be created. The opportunities will raise more qualitative issues, including both the organizational structures and the instructional processes in higher education, since they can be affected by these developments.

Harvard University

Submitted July 1973

Revised November 1973

REFERENCES

- [1] Wyatt, Joe B., "The Harvard Plan," *Networks for Research and Education: Sharing Computer and Information Resources Nationwide*. Eds. Greenberger, Aronfsky, McKenney and Massy. MIT Press, 1973. Chapter 33.
- [2] Sharpe, William F., *The Economics of Computers*. Columbia University Press, 1969. Chapter 9.
- [3] Watson, Thomas J., Jr., Keynote address to SHARE meeting, Fall 1965. San Diego, California.
- [4] Polley, Vincent L., *Computing at Harvard*, Office for Information Technology Survey, April 1973.
- [5] Stuehrk, Carl F., "Digital Data Service," Proceedings of the Fall EDUCOM Conference, October 1973.
- [6] Talbert, Lee R., "Before the Federal Communications Commission in the Matter of Packet Communications, Inc. Application for authority to institute and operate a packet-switched communications network in the Continental United States," January 23, 1973.
- [7] Holt, Charles C., "A System of Information Centers for Research and Decision Making," *The American Economic Review*, May 1970.
- [8] *Networks for Research and Education: Sharing Computer and Information Resources Nationwide*, Greenberger, Aronfsky, McKenney and Massy. MIT Press, 1973.