RJH

BUSINESS AND SCIENTIFIC COMPUTING - A PEACEFUL COEXISTENCE

Robert J. Harvey
Mathematics and Computing

Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

October 18, 1974

I first would like to point out to you that I have altered the title of my presentation. I have taken the liberty to insert the word "peaceful". One definition of the word coexistence that I found was the following: Simultaneous existence through a policy of mutual non-interference, usually referring to two nations differing widely in ideology. Although this may have been the situation some time ago with respect to business and scientific computing, I don't believe things are all that bad today.

Secondly, I will approach the subject of coexistence assuming that we are thinking of a situation where both business and scientific computing are sharing the same computer hardware. Under these conditions it certainly is better if the coexistence is a peaceful one.

Now let me outline what I want to say to you during the next few minutes. I will begin by pointing out a few of the well known areas where business and scientific computing have conflicted with one another. I am sure you will all recognize the situations as I describe them to you, however, you probably will not all see these situations in the same way. That is, some of us will recognize them from the viewpoint of business data processing, while others of us will see them from the scientific computing point of view. Once we all remember some of the areas that have caused business and scientific computing to disagree, I will introduce what might

become a new, aggravating situation as far as peaceful coexistence is concerned. Finally, I will point out what I consider to be some promising developments and conclude by submitting for your consideration what I believe to be the key to the realization of peaceful coexistence.

Areas of Conflict:

Now let's remember some of the areas of conflict between the two groups.

1. Peak Loads: While the job load on a strictly scientific computer center is more or less uniform in time, this is not at all the case for the job load initiated by the business side of an organization. The business computing job load is full of peaks. There is the end of the month peak made up of payroll time card accounting, pay checks, general ledger and countless monthly reports. Then there are the peak work loads that occur quarterly, semi-annually, and let's not forget the end of the fiscal year. Hundreds of data files that must be updated, merged, sorted, averaged and reports generated; all within a short period of time. Each peak work load is made up of the business and accounting transactions that have occurred to date.

In addition to the predicted peak loads that occur in business computing there are always those unscheduled, unpredicted demands for particular runs on the computer to prepare special, one-time reports. These requests usually must be satisfied in a short time period and so they produce another peak in the computer job load.

It sometimes appears that the business data processing machines are used just to cover peak demands. After all, one could not get all of the accounting done nowadays without the use of a computing machine and the work load occurs in peaks.

Now I am not trying to say that peak demands never occur in scienti-

fic computing; to be sure, they do. Scientists do sometimes work up to the last minute before dashing off to present a result at a meeting or conference. However, the peak job loads do not appear with any sort of regularity and certainly they don't seem to have the life or death completion dates associated with them.

Sharp peaks in a job load usually cause scheduling and priority problems and conflicts are bound to develop.

2. Security and Data Integrity: All kinds of confidential data are handled by the business computing group. These confidential data are made up of personnel files, payrolls, and other corporate data covering a variety of sensitive levels. The employee supplies personnel information to his employer assuming it will be handled in a strictly confidential manner. Payroll data must be handled in a secure environment since widespread knowledge of salary data is not a common practice. The revelation of certain corporate data could cause a company to lose its position to the competition. This situation is not as applicable to those of us present, however there do exist data that are handled on a need to know basis.

With the exception of centers that handle classified data, scientific computing is not, by and large, concerned with security.

Data Integrity

Naturally, a scientist does not like to lose several days work as a result of a misplaced or overwritten magnetic tape, and the operations group must exercise care and caution to avoid such incidents. While costly and valuable scientific data are treated properly, a great deal of the work done on the computer is made up of short, one shot jobs that cost less to reconstruct than it would to provide means to ensure data integrity. On the business computing side, there are elaborate steps taken at all times to ensure data integrity. These two topics are in some sense

part of a larger area that I would like to mention next.

3. Operational Conflicts. There are many items in this category, but I will mention only a couple of them here. In business computing there seems to be a large number of special forms. When this requirement is intermixed with scientific computing, there results a constant changing of paper and special forms on the line printers.

∴ FOG is BDP

A second operational procedure common to business computing is the job to job dependency. One can not run job B until job A has been successfully completed because the output from job A is used as input to job B. The situation is really worse than that because there is a string of several of these jobs that must be completed by 8 A.M. Everyone knows that something always goes wrong during such a process and when that happens, one can not just set the whole job aside until someone comes in during normal working hours. The solution that has been implemented in some cases is to have specially trained expeditors on hand to supervise the operation during such runs. Naturally these specialists are not the everyday (or night) computer operators. So within no time at all, there is a war in the computer room. I will mention later a similar situation that occurred on the scientific side, but now let us continue and explore some other areas of possible conflict between business and scientific computing.

4. Software Operating Systems. I understand that most business computer people are strong believers in labeled tapes. That is, each reel of magnetic tape must contain as its first record, a block of data that identifies that particular reel from any other reel of tape. The software system provides the user with a means of utilizing that identification block for a variety of purposes. For example, he may assure himself that his program has access to the correct magnetic tape and hence avoid possible loss of time and money. He may also use the identification record to

establish when the tape was last used, how many times it has been used, etc., etc. Of course the system itself may use the tape lable to determine if the person requesting the tape is in fact authorized to make use of that particular data, and so on.

Once again, the scientific user does not want to have his data mistakenly overwritten, but by and large he is unwilling to put up with the inconveniences that such interlock procedures imply.

Also, there is the whole area of binary coded decimal versus binary. Some people, so I am told, want to do arithmetic using BCD data.

Above all, business data processing users demand system stability. When the software system is down for a long period of time, they are not very impressed by the promise that when things get back to normal, the new system will run the job 10% faster, or that the old bug (that by now everyone has coded around) will be removed. The scientific user is not overjoyed when the system crashes but he does want new features made available to him.

The business computer users do rely on the system to be consistent over a long period of time. The payroll job must run correctly month after month, without large investments in program maintenance due to changes in the software operating sytems.

5. Hardware Needs: The scientific computer user community wants computer main frames with fast floating-point arithmetic, and many significant digits. The business user wants plenty of magnetic tape drives or removable disc pacs for his sort/merge jobs, and field manipulating capabilities in the main frame hardware.

So far, I have tried to mention some of the differences of opinion that have been the basis for the conflict between business and scientific computer users over the past years. I have made no attempt to list all

the differences, but rather to recall some of them so that we all get in the proper frame of mind.

Now let me introduce a new area that could jeopardize the peaceful coexistence of these two groups. This topic has been vigorously debated over the past few months. I personally have not heard this subject mentioned as a possible threat to peaceful coexistence and it may be that it will not develop that way. The area I have in mind is that of management controls over computer resources. I am sure you are all aware of the subject. Some spokesmen from scientific computer centers have taken the stand that although controls of this kind serve a useful and necessary function, some forms of formal management controls and procedures are in fact counter-productive to using the computer as a tool for scientific research. In the business community such formal control practices have long been accepted and adhered to. It is the hope of some of us from the scientific side that everyone will come to realize that although we may be able to coexist peacefully with business computing, we are different. We are different because we don't always know what we are doing, and if we did, it would no longer be research. Therefore, I don't believe that just because we use the same equipment that we should also necessarily use the same management control procedures.

Changing Situation:

Attitudes and techniques have changed over the years. I had the opportunity to spend some time during the late 1950's and early 1960's with a group of physicists at Berkeley who were engaged in experimental

high energy physics research. Although some of us didn't pay much attention to it, many of the problems we were trying to solve and a good deal of our concepts and attitudes towards computing were similar to those of people engaged in business data processing. We struggled to run inter-related job streams. We demanded block time on the computers, and we had our own highly skilled computer expeditors who would take over the Computer Center to run the jobs using eight tape drives. We tried to use an operating system on the IBM 709 called S.O.S. and were frustrated when the on-line printer clanked out: "IF THIS IS A GO-JOB, IT IS NO-GO; IF THIS IS A NO-GO JOB, IT IS COMPLETED". We also made peak load demands on the Computer Center. There were times when peaceful coexistence was questionable within the scientific computing community. Since then, both the computer hardware and software systems have improved. I believe that the introduction of multiprogramming forced us to re-think some of our methods and many of the old problems went away.

A concept that has become increasingly important to computer people is the concept of a computerized data base. Of course, the business users have been dealing with problems involving data bases all along. The current emphasis on research in the field of energy and environment, for example, means that more scientific computer users are thinking about these same problems. Many of them even know what a fully-inverted file is. So here we have an area of mutual interest that can be used as help to develop peaceful coexistence.

I have pointed out some of the problems that have arisen when business and scientific computing is done on the same machine. I have pointed to one instance where some strongly maintained ideas in the scientific community have changed. I also gave an example of business

and scientific computer users becoming more interested in similar problems.

The question of peaceful coexistence has no simple answer. The approach is probably different depending on the corporate mission and the size of the computer installation. Not everyone has the same sorts of resources. Not everyone agrees on the seriousness of the areas of conflict that I just mentioned.

Like many considerations, I believe the key to a realization of the peaceful coexistence of scientific and business computing is people. Before the technical problems can be solved, everyone must agree that peaceful coexistence is what they want to achieve. By everyone, I mean not only the Computer Center Management, but representatives of each area that will place demands on the computer.

I think it is possible to achieve peaceful coexistence if the people involved get together and make sure that everyone understands the ground rules and agrees to stick to those rules. I don't want to imply that suddenly all the problems will vanish. However, people will know what to expect and I think that is important.

Now I am not advocating that everyone should rush out and put their business and scientific computing on the same machine. That may or may not be the thing to do. However, for those who want to try, I leave the following message: Make certain everyone is on board and hammer out the ground rules and I think you will maximize the probability of success.