SDS

MEMORANDUM

TO Distribution

DATE March 31, 1969

FROM George Boyd

PP-GB-1211

SUBJECT: The Future of the Time-Sharing Market and the Product Implications for the Sigma Series

REFERENCE

This memo represents the Product Planning position on the time-sharing market which will exist or be coming into existence in the early 1970s, and defines the software products that will be required to cover that market. The final product decisions will be based on this document and the modifying information received in response to it from Marketing and Programming; formal responses are hereby solicited. The final question will of course be one of what products SDS should build for the market, but since there will be further discussion and compromise before that point is reached it is requested that Marketing's response include an evaluation of the market definition contained here, and that Programming respond to the practicality of the development projects described here. Formal written reponses are requested by April 7, 1969. The market percentages and other numbers used in this document have all been extrapolated by the author from small samples and from widely varying opinions of the experts in the field, and as such can only be considered as illustrative. The probable variations in the data are, however, not considered great enough to change the results of this study, so the involved numbers will be quoted without further qualification.

It is quite clear that the time-sharing market is going to continue growing at a rapid pace until more than one-half of the dollar volume of computers sold in the early 1970s will be sold into environments in which time-sharing is an important consideration. Time-sharing as referenced above includes all remote time-shared utilization of computers, and consists of both remote batch and conversational usage. The time-sharing usage of computers can be broken down into three conveniently different groups; remote batch, on-line non-scientific, and scientific problem solving. Eventually the time-sharing market will be divided into 50 percent remote batch, 35 percent on-line non-scientific and 15 percent scientific and problem solving applications. The market should not reach that point until the middle or late 1970s though, and the ratios will be significantly different over the life of the currently planned Sigma machines, or up through 1974. The current percentage of the time-sharing market that is remote batch is about 50 percent, which should not change significantly as a percentage through 1974. Time-sharing in non-scientific applications amounts to less than 10 percent of the current market but should increase to 25 percent by 1974. Scientific time-sharing and on-line scientific problem solving is currently about 40 percent of the market and should decrease to 25 percent by 1974 as on-line commerical and remote batch usage increase at a greater relative rate. The market in which time-sharing is at least an important consideration should amount to about one-half of the total computer market by 1974.

There are a large number of basic questions about the shape and the future of the time-sharing market at the current time, and about how SDS should prepare to meet the demands of that market. The picture is really quite complicated, primarily because the time-sharing market is still very new. The initial thrust into this market was by entreprenuers who were willing to risk the changing environment, and computer companies who essentially inherited the market by default. The market was initially shaped by the available machines and available software rather than by customer pressures. The time-sharing market is now about to enter its second stage of growth on the way to legitimacy, with established service organizations and computer companies entering the market in a serious way. The market is also expanding to include a broader segment of the population of computer users. The initial customer support for time-sharing came from scientific users who were solving short term individual problems, and people who were simply interested in time-sharing. With the availability of more capable and reliable hardware and more elaborate software, longer term and more serious applications will be put on time-sharing systems, with non-scientific and BDP applications being the most important candidates. It is estimated that 85 percent of time-sharing resources in the future will be involved with non-scientific or BDP applications including remote batch. The market will also break down into several different categories of customers; the service bureau, the time-sharing utility, the in-house batch system with time-sharing, the in-house time-sharing system with dedicated applications, etc.

The products that SDS must supply to capture a large share of the future time-sharing market can only be determined if a good model of the future market can be built, and it is with the idea in mind of building that model that the following discussions are forwarded.

Non-Scientific Applications

Non-scientific time-sharing applications will include on-line and remote batch use of computers for BDP or business oriented functions, such as accounting, inquiry, data gathering, ticket systems, manufacturing control, distribution system control etc. The non-scientific portion of the time-sharing market will be divided into two thirds remote batch and one third on-line utilization by 1974. Non-scientific and BDP time-sharing will make up 75 percent of the time-sharing market and about 40 percent of the total computer market.

Non-scientific and BDP applications are characterized by quite different physical attributes than scientific time-sharing. The major differences are associated with far higher data input rates, more frequent interactions, less computation but more I/O per interaction, and manipulation of large numbers of small pieces of data within large data bases. In addition, a non-scientific or BDP application package will generally give rise to a significant amount of batch work. For example, the BMA system in Salt Lake City generally has to bill on-line users for an equivalent amount for both batch and on-line work, and that one for one ratio should remain typical. The type of batch work generated involves sorts, merges, and report generation on a periodic basis. The majority of on-line commercial packages are also used primarily during daylight hours and on week days only, so that a well loaded general purpose non-scientific time-sharing utility system will do well to bill more than one quarter of the machine resources for direct on-line usage, and an additional one quarter for batch support of the on-line users, although these numbers can vary widely. Most non-scientific time-sharing service bureaus will use batch processing to smooth the system loading when the on-line users are active and to fill the idle time of the machine. The same should also be true of in-house GP systems. That fact is important but not critical because of the fact that the billable amount of hardware in a large time-sharing system is fairly small, and because the hardware costs are not an extremely critical part of the total cost of operating a service bureau. Another important point is that in any large scale time-sharing service bureau there will be one or two customers who will utilize the system all hours of the day and night, so that the batch and on-line operations will have to run concurrently a significant part of the time, and it will not be possible to switch monitor systems in order to do batch work. Utilization of the peripherals to the full also means running line printers and card equipment all of the time that the system is up, since a large Sigma 7 time-sharing system in a commercial time-sharing environment, for example, would generate enough output to keep an average of 2 or 3 high-speed line printers busy full time. In addition, remote batch will probably account for half of the total billings in a typical non-scientific time-sharing service bureau by 1973.

This market will be covered primarily by the existing batch oriented service bureaus, 80 percent of whom now feel that they will be providing time-sharing services by 1973. The existing BDP oriented service bureaus will probably manage to take virtually the entire non-scientific time-sharing market because of the fact that it is basically a people and support business rather than a service business like scientific time-sharing. Some of the new time-sharing service bureaus will break into the non-scientific time-sharing market, but they will do it by offering special terminal services in the form of proprietary and very special purpose applications programs rather than general BDP service with all of the support that is implied.

The non-scientific time-sharing market will require software that differs significantly from any that SDS has under development at the current time. The monitor services required for file retrieval and dealing with large files in a fast and efficient manner are well in excess of those available under UTS, even in the new file management system. In addition, the monitor coding to back up the user services would have to be considerably more capable of using large numbers of large and diverse secondary storage devices in an efficient manner. The time-sharing portion of the system would have to be constructed to work well in the face of the high level of requests for interaction that a BDP oriented time-sharing system must face. The monitor would have to be able to work on a transaction basis with automatic file update, automatic data replication, and probably automatic partial promotion to more accessible devices of in-use data, in order to provide the type of safety features and efficiency that the BDP environment will require. In short, the monitor services would have to be more extensive than those of UTS, with a specific orientation toward the types of files that non-scientific and BDP applications packages generally have to deal with. The monitor would also have to provide an efficient general batch capability, which would definitely have to include a multi-programming facility.

The basic non-scientific application program categories include accounting, manufacturing control, inquiry, data collection, management information, and administrative systems. While a larger number of categories could easily be defined, the categories of accounting and inquiry/update will still amount to the majority of non-scientific time-sharing applications. The non-scientific load will eventually be split in a fairly even way between remote batch and on-line terminal applications, where time-sharing is involved.

The requirements for software to support a non-scientific or BDP oriented time-sharing system will be very much the same, whether for a service bureau or for in-house usage. As a matter of fact, the in-house requirements as an aggregate may be more demanding than those of the service bureaus, since a service bureau can selectively limit its market.

Scientific and Engineering Market

Scientific and engineering applications programs and problem solving will make up approximately 20 to 25 percent of the time-sharing market through 1974. The types of services that will be required by the user in this portion of the market will not differ substantially from those that are currently being made available, except that fail soft will become the general rule. The major difference in the market will be in that there will be a much wider spectrum of types of machines and types of software to choose from, and that time-sharing services will be offered by a broader class of suppliers. The machines that were used initially in this market had little ability to provide batch services, either to provide support for the on-line users or to gain additional income by using idle time and off hours for batch work. The major element that kept batch processing from entering the picture was the inability of the systems to do very much batch work even if the hardware and software had allowed concurrent batch and on-line operation; the initially using machines were all rather small and slow by current standards. With the current generation of machines, the scientific time-sharing user will be able to utilize a portion of an in-house batch system as well as buy terminal time from a service bureau or time-sharing utility.

The major motivation in a time-sharing utility environment will of course be to make money. The secondary but still very important point in a utility business will involve operational considerations. In the current market, hardware costs generally run from 10 to 25 percent of the billings for a fully loaded system.

Of that hardware associated portion of the billings, a substantial fraction is proportional to the number of on-line users that are being supported, as are almost all of the other costs of running a service bureau. So, the efficiency of utilization of the hardware becomes fairly unimportant on a general base, and a utility service might very well choose to suffer the idle time of the system rather than try to broaden the base of their service to fill the idle time with batch processing, which is a very different kind of business for them. On the other hand, service bureaus that are providing non-scientific time-sharing services and batch services concurrently should find it very easy to provide scientific time-sharing also. The possible inefficiency of their attempts to do so would be relatively unimportant because of the fact that the majority of the machine expense will be underwritten by non-scientific and BDP work, and providing scientific time-sharing will amount to selling computer time only and will provide a particularly painless income since little support would be required.

The general conclusion is then that there will be scientific and application time-sharing utility services in existence in the future, and that those businesses can get by without providing sophisticated batch services. The utilities will not, however, have the complete market to themselves. The large commercial service bureaus will take a large portion of the market. In addition, an important part of the utility market will continue to use small machines and provide service to a relatively small number of users, since the threshold price for a machine in this market should continue to be very important, as should the expansion price.

Market Summary

The time-sharing market in the future will involve enough differences in kind of service and a large enough amount of overhead in operation that the hardware cost per line will not be the absolute driving competitive factor. Specifically, there will be budget limited and special purpose parts of the market that will be significantly large despite the relatively high hardware costs per line that they will have to bear. While it is safe to say that the majority of time-sharing business will be involved with broadly based systems supporting both batch and time-sharing operations on the same equipment, time-sharing will definitely be saleable on machines which are either not optimally sized for efficient utilization or which do not support general batch operation. Most of those minor market machine sales will be made on the basis of low price threshold or the existence of specialized software for a particular application or type of application.

The majority of the market will be associated with support of remote batch and non-scientific on-line applications, which are tied together in their basic system requirements to a significant degree, and which will amount to 75 percent of the time-sharing market by 1974. The remainder of the market will be problem solving and scientific applications. The largest part of the market will be covered by the same software for both in-house and service bureau usage, with the aggregate set of in-house software requirements being at least as stringent as the service bureau requirements. On the other hand, no single system will be able to satisfy the overall requirements of the entire market.

The SDS share of the time-sharing market will depend very heavily on the degree of penetration of the BDP market by the company. The non-scientific and remote batch portions of the time-sharing market will be so closely tied into BDP batch that the batch facility will be the driving function. Only a small portion of the remote batch and commercial on-line market would be available to SDS based on current BDP capability, and that portion would be associated with scientific shops for remote batch, and special non-scientific or BDP applications for on-line utilization. The scientific and problem solving market will be split between in-house batch systems with some on-line scientific work for problem solving plus program development, service bureaus offering scientific time-sharing in addition to BDP batch and commercial on-line, and systems dedicated to scientific problem solving and special applications.

Product Requirements

The currently defined SDS software and hardware will cover the majority of the time-sharing market through 1974, but additional products will be required. BTM and UTS on Sigma 5 and 7 will be adequate for approximately one half of the time-sharing market, in terms of dollar value of equipment, but will leave three major market areas of significant size for which other products should be defined. The three areas are characterized by, 1) a market with a lower price threshold than is possible with BTM on a Sigma 5, 2) a market with higher performance requirements than can be satisfied by UTS on a Sigma 7, and 3) a market requiring significant differences in kind of service provided as well as the number of users supported in a time-sharing environment. The products that must be defined to cover these portions of the market are described in the following paragraphs, and include, 1) a Sigma 3 time-sharing system, 2) an extended UTS system for Sigma 9, and 3) MPM, a Sigma 9 time-sharing system. It should be pointed out that having a market covered with software and hardware does not necessarily mean that SDS will be able to make a significant penetration of that market. In fact, it is in the largest part of the future market, the BDP time-sharing portion, that SDS will be least competitive.

The physical products that are required for the future, or which would show an acceptable return on investment, are as follows. First, a Sigma 3 time-sharing system appears justifiable on the basis of the fact that we should be able to capture a significant portion of a reasonably large market with such a system. A software system built on RBM and able to support 32 Basic or Fortran users in dedicated systems or 16 general purpose users would be the goal. The BTM will provide continuity in price and performance from the Sigma 3 System up to UTS on a Sigma 7. Sigma 9 will have to be provided with a version of UTS expanded to take advantage of the fail soft, reconfiguration and multi-processor capabilities of the hardware, and extended to provide a broader base of batch services for use in the non-scientific and business data processing environment. The MPM product is required to cover the time-sharing utility and special application market. The MPM system will have to provide support for scientific and engineering applications programs, scientific problem solving, and special purpose commercial data processing applications programs. A batch capability is of interest under MPM, but only to the extent necessary to support the on-line applications packages. The emphasis of MPM should be to provide a general type of service for large scale time-sharing applications programs, and it should not be unduly complicated in order to provide a broader base of services than is necessary for its market. The basic premise behind MPM should be that it is designed to support a subset of all possible applications programs in an efficient and intelligent manner, and that it perform that task well at the cost of generality. That position will produce a viable product since there will be many applications programs that general purpose time-sharing systems will not be able to support appropriately that an MPM system will support, and since the time-sharing utility companies will be able to offer specific proprietary software packages which will not be available elsewhere or need auxilliary support. MPM should also serve as a good base for dedicated application time-sharing systems, such as ticket reservation or data gathering systems.

List from someone please!!

Sigma 3 Time-Sharing System

The Sigma 3 time-sharing system will cover 5 percent of the total time-sharing market, and will contribute over $40 million in sales by selling 150 Sigma 3 systems with an average price of $300K and a minimum price of $250K. The time-sharing system will have to provide support for 16 general purpose users or 32 users in a dedicated Basic or Fortran environment. The software will be built on RBM, and will be constructed using variable partitions for programs in the manner of UTS, rather than a single partition like BTM, but would of course be much simpler than UTS. This approach is practical for Sigma 3 because it has a base and index register structure which will permit the dynamic relocation of programs required by the variable partition approach. The development cost of the system would be approximately 100 man months, and the system would have to be available before the end of 1970. A disc pack or RAD providing approximately 15 million bytes of storage, at a price of less than $50K, is a necessary element of the Sigma 3 system which is not really feasible without it.

Sigma 9 UTS

The Sigma 9 UTS system will be based on the Sigma 7 multi-programming UTS design, and will be compatible with the Sigma 7 product. The total software system on Sigma 9 will consist of the existing processors and subsystems and a monitor which is a compatible superset of UTM. The monitor would be reimplemented to take advantage of the fail soft, larger memory, dynamic reconfiguration, and shared core multi-processor facilities of the Sigma 9, and would be extended to provide a broader base of services in the areas of file management and communications. Compatibility will be limited to the Meta-Symbol source level in order to allow some modification of the monitor service call formats where that is desirable, but no other incompatibilities should be necessary. This approach allows existing programs and processors to be converted without modification, but does require reassembly of all Meta-Symbol written programs.

The above position on Sigma 9 UTS is the appropriate goal for the main Sigma 9 software. It is clearly the case, though, that there is some risk in trying to get to a fully expanded Sigma 9 UTS system from the Sigma 7 UTS, and the feasibility of achieving a multi-processing system on that basis must of course be investigated before any commitments can be made to that development. However, another way of viewing the Sigma 9 UTS development would be to consider the implementation from scratch of a new monitor which was a completely compatible superset of Sigma 7 UTS, or compatible at all levels other than the object language level as a possible second choice. Somewhere between the two extremes of 1) completely reimplementing UTS for Sigma 9 in a Sigma 7 compatible manner and 2) changing a minimum of the existing Sigma 7 UTS code to produce a Sigma 9 UTS, there should be a method of building the desired product. There is of course the possiblity that the multi-processor extensions of Sigma 7 UTS to Sigma 9 are just not reasonable and that they ought not to be done, but that is yet to be determined, and must be studied. Completely reimplementing the involved portions of the Sigma 7 monitor for Sigma 9 UTS would require approximately 35 man years of development (70K words of code), while the simplest conceivable adaptation of the existing code would require 5 man years.

The Sigma 9 UTS will provide continuity with the BPM and UTS systems on Sigma 5 and 7, and will cover the general purpose scientific market very well while providing coverage of the non-scientific and BDP sale of Sigma 9 systems as well. Continuity will be very important to the sale of Sigma 9 systems, as well as to the sale of Sigma 7s. The extended UTS system for Sigma 9 should cover approximately 20 percent of the time-sharing market for SDS, while providing more competitive coverage of a portion of what is currently the Sigma 7 market, and thereby allowing an increment in Sigma 7 sales by providing a compatible upgrade. The Sigma 9 UTS would be available to the field in the middle of 1971 at a development cost of 15 man years, and would be responsible for the sale of 100 Sigma 9 systems.

<u>Shared File Sigma 9 UTS System</u>

The question of building a shared file multi-computer system for Sigma 9 has very little to do with the question of attempting to develop a shared core multi-processor system. The value of the two approaches are very different, and have application to different types of uses, and present very different cost and fail soft cycles. The shared core multi-processor provides a smooth growth of a fail soft system from a small non-redundant configuration up through partial redundancy to a dual CPU fail soft configuration. With a shared file system, redundancy is possible in the same manner from a small configuration to the same point of single CPU redundancy as a shared core system, but at that point full fail soft involves replication of the entire configuration rather than adding a CPU. If a CPU goes down in a shared file system one full complement of equipment is lost. While many people who buy more than one Sigma 9 might like to have shared file software, few people who are concerned primarily with fail soft or system balance will accept the shared file approach. Where CPU horse power is concerned, shared core multi-processing is the best solution by far, in most cases and from most viewpoints.

A shared file system should be implemented for Sigma 9 UTS, but a shared core multiprocessor UTS must also be built if it is possible.

Sigma 9 MPM

The Sigma 9 MPM system will be the second generation of the 940 system, and should cover quite a bit more than the traditional 940 market. The software will be much more substantial than the 940 software because of the more capable hardware, because of the experience that has been gained with time-sharing software, and because of the increased sophistication of the customers and increasing competition for those customers. The design emphasis of the system will be put on support of applications packages and scientific problem solving, but batch and remote batch capabilities will be built into the system for support of the on-line users. The advantage of the MPM system over UTS will be in its specialized orientation toward the terminal oriented market, and will consist of greater efficiency for time-sharing users and simpler to use but more capable on-line services for the user. The implementation emphasis will be placed on fail soft, with dynamic reconfiguration and multi-processing being included also. The system design objective will be to support 200 users on a fail soft configuration with batch support for the on-line users. The system will appear similar to the CSC produced TSU specification, except with improved batch and remote batch capabilities, greater emphasis on fail soft, multiprocessing, and a broader base of support for application systems.

There is no extreme requirement for compatibility with any existing SDS software which would be imposed on the MPM system for the market that has been defined here. If MPM is to be of value as a backup or parallel for Sigma 9 UTS, however, it must be compatible with the Sigma 7 UTS system and with BPM. There is some risk involved in producing the Sigma 9 UTS system, and a backup in the form of MPM is of interest. Whether or not the question of the role of backup is important enough to be worth imposing major compatibility constraints on MPM, the system should still attempt to be as compatible as possible at all levels where the concession to compatibility is not an important liability.

The MPM system will uniquely cover 10 percent of the time-sharing market for SDS and will cover an additional segment of the market in a more competitive manner than other SDS products. SDS can expect to sell 50 Sigma 9 MPM systems into the time-sharing utility and special applications market, and 15 machines for more general purpose use, at an average price of $2 million.

The development of the MPM system will require 75 man years, plus development of the required subsystems and processors. The complete system can be scheduled for initial delivery to the field in the fourth quarter of 1971.

GCB

George Boyd

GB:cg

COPIES TO:

A. Bongarzone
W. Colby
D. Cota
P. England
W. Gable
P. Likins
E. Maso
D. McGurk
J. Mendelson
M. Palevsky
L. Perillo
W. Shultz
R. Spinrad
J. Stoffer