

**ASDD SAN JOSE LABORATORY**

**August 23, 1963**

**MEMORANDUM TO: E.R. Piere, CHQ, T.J. Watson Research Center  
B.O. Evans, DSD Hq., Poughkeepsie**

**SUBJECT: The AEC Computer Situation**

**As a result of my recent tour of duty as a consultant at Los Alamos and conversations with many people, both within the AEC and IBM, I became concerned about IBM's position and trend with respect to the "Gold Chip" customers. The enclosed paper on the subject gives the picture as it appears from my particular vantage point.**

**The intent is to be helpful, not critical of any individuals or their efforts. The opinions expressed are my own and do not represent an official ASDD position.**

*H. G. Kolsky*

**H. G. Kolsky** ✓

**dmc**

**Encl.**

**cc: R. B. Johnson**

IBM CONFIDENTIAL

SOME COMMENTS ON THE AEC COMPUTER SITUATION

August 20, 1963

Dr. Harwood Kolsky  
San Jose ASDD Laboratory

I The Problem

During the past few months there has been an increasing amount of concern expressed within IBM concerning our status with respect to the Atomic Energy Commission and its Laboratories, particularly with regard to their future computers. There have been a number of meetings held in recent weeks within DP and DS on this subject.

Without going into detail, the cold facts show that almost all of the AEC Laboratories presently have competitive equipment installed or on order; much of it earmarked to replace IBM equipment. Livermore presently has on order a CDC 6600, which is a STRETCH-class machine or better. There also have been meetings within the AEC computing community at which a strong case has been made for all the Laboratories to standardize on CDC equipment (the 3600 in particular). The arguments

were mainly that this equipment is more modern and more nearly fits the standard scientific problems of the AEC than does the IBM product line which they feel is now becoming too expensive and obsolete technologically. The on-order position for the AEC at present shows that IBM has something like 10% of the machines on-order.

## II IBM-AEC History

If one contrasts this with the historical AEC position, the trend is indeed frightening. Ever since World War II the AEC has been a leader in scientific computation. At Los Alamos, for example, IBM accounting machines were brought in and used on lengthy scientific problems at a time (1944-45) when this was a highly unusual procedure.

The AEC Laboratories have traditionally used priorities to obtain the first of IBM's new scientific machines off the line; the CPC Model 1 and 2, 701 and 704, etc. They have been real leaders in the scientific computation field in this country and have always, at least until recently, shown a preference for IBM equipment.

In addition to serving as a pace-setter in the encouragement of commercial equipment development, the AEC Laboratories have also been active in designing and building their own machines. The original Van Neumann "Maniac" type machines were heavily subsidized by the

AEC Laboratories. Argonne, Oak Ridge and Los Alamos were three that built variations on the original Maniac design. These projects certainly encouraged commercial entry into the field. The IBM 701 was directly influenced by them.

Lawrence Radiation Laboratory of Livermore, when it ordered the LARC computer from Sperry-Rand certainly pushed the state-of-the-art forward, and was very instrumental in encouraging IBM to undertake the STRETCH contract with Los Alamos a year later. (Livermore later also bought a STRETCH).

Clearly, the present concern over the AEC situation is very well founded since it represents an important area in which IBM has been an undisputed leader for many years, but which is now deserting us, it seems. What is more, there is no obvious counter-action on the horizon likely to reverse this trend. One should ask first what has gone wrong and secondly, what, if anything, can be done about it.

One point of view within IBM considers the AEC problem as simply a marketing problem. This point of view says that all we need do is apply the proper marketing posture and we can save the day the same way Mr. Learson saved the 702 market ten years ago.

The situation is certainly not this simple. One of the characteristics of these "Gold Chip" customers (which include perhaps three dozen organizations, including the AEC Laboratories), is the fact that

they are extremely competent in their own fields and very knowledgeable in the field of computers. A typical marketing approach is rather ineffective on these people. They are much more apt to be convinced by a straightforward scientific discussion of the technical problems and solutions. One does not need a massive marketing organization to deal with the Gold Chip customers, simply because there are so few of them. A small handful of very competent scientifically oriented people can make the rounds of all three dozen or so of these customers without having a large marketing backup. (Incidentally, this is exactly the type of strategy which CDC has used with devastating results.)

### III Some Comments Heard at Los Alamos

I recently returned from spending some time as a technical consultant for the AEC in Los Alamos. Although my duties there have no direct bearing upon computers, and certainly not upon the computer requirements generation; nevertheless, I could scarcely avoid having informal discussions with my old friends there who are vitally concerned with the use of computers in solving their problems. After a few days I began to appreciate their viewpoint again, which looks quite different from the one seen from the IBM side of the table.

The following comments were typical: "What's wrong with IBM anyway; is it coming apart at the seams?" "Did you see that GE 225 in the machine room?" (I had...it replaced two IBM 1401's.) "What's wrong with IBM; don't you have anything to match competition anymore?" "What's happening to IBM; aren't you going to bring out any new computers?" "You're leaving us no choice but to go to the 6600." "We've always have very good luck with our IBM machines and we feel as though we understand you and you understand us; but what is wrong with you now?" "IBM stock has been going down and CDC has been going up; does this mean anything?"

Other questions were raised as a result of recent personnel changes that they heard about. "Is there some kind of big shakeup within the Company?" "What's the real reason behind...etc?"

In July there was a presentation by DSD managers to a group of LASL scientific leaders concerning IBM's future plans, including some about NPL. I later talked individually to most of the LASL people who attended this presentation. The almost universal reaction was: "It sounds as if IBM is not interested in our business anymore." I asked specifically how they could ever get such an erroneous idea as that. Their answers could point to nothing specific which was said in the presentation, but nevertheless, they

came away with the impression that IBM was really not interested in building a machine above the STRETCH range, nor in scientific computation in general. NPL, as presented, did not thrill them. I heard comments such as: "Well, IBM is willing to take our money if we want something special, but they don't particularly care to do it." One highly placed person told me: "It sounded to me as though we were being invited to go to the 6600." Another Group Leader said: "It looks as though we will just have to leave IBM out this next generation of computers." Another said: "As a result we are re-considering the possibility of building Maniac III ourselves." There were also the usual comments of praise about CDC and the impressive presentations they had made.

#### IV What Should Be Done About It?

The question is not so much how did this situation arise, but what should IBM do now to answer the crisis? My thesis is that only a bold stroke will save the day... in particular, the stroke of signing a contract with one of the AEC Laboratories for a Super-STRETCH machine (meaning one at least five times STRETCH performance on AEC problems).

There are several arguments supporting this contention:

1. The competitive situation mentioned above in which the AEC Laboratories are abandoning IBM equipment, at least in their future plans.
2. IBM presently has in hand superior high speed circuits (Impact), very high speed small memories, and the large bulk memories which provide the building blocks for an extremely high speed machine while still retaining a relatively simple CPU design. In fact, relatively, we are much better off now than we were at the start of the STRETCH project.
3. Because of the marketing plans for NPL there is no real internal incentive for driving toward a very high speed computer unless it is done on contract, especially oriented toward our "Gold Chip" customers.
4. The recent granting of the Redman contract by one such customer is an extremely important point here. Many preliminary data flow and logical interconnection problems for the high speed components will be solved as part of this relatively small program. Much of this experience could be directly applied towards the designing of a larger CPU.
5. Indications are that the timing is ripe within the AEC. While LRL has a CDC 6600 on order, Los Alamos does not yet have a comparable machine on order. Although they are presently negotiating for a 7094 Model 2, which is in the STRETCH speed range, it is not considered to be a STRETCH successor. LASL is certainly feeling the rivalry of LRL and NYU in their search for larger computers.



There are a number of arguments against undertaking such a contract. The following comments attempt to answer some of these.

V Fallacies of the "Market Pyramid"

The computing market has often been described as a pyramid in which a few extremely large users, including the AEC Laboratories, are at the top of the pyramid. The market broadens down to include aerospace industries and other government supported groups next, then the larger commercial users and finally, layers of smaller and smaller commercial users. The number of installations increases rapidly as one goes down in size.

If one optimizes a computer design on the basis of income, it always turns out that a design aimed for the middle of the pyramid or the lower-middle of the pyramid will result in far greater total income than one tailored either for the top, "Gold Chip" customers, or for the bottom myriad of small potential users where the competition is quite fierce.

The NPL market philosophy has taken exactly this point of view and seems to have optimized their whole compatible structure around total profit, which means they have been optimizing around the middle of the market pyramid.

Unfortunately, profit is not the only thing which is important in this game. Like any other field of endeavor, there are a lot of subjective, philosophical points that are hidden from these hard market figures. There is a great deal of "fashion setting," "pace-setting," and "trail-breaking" philosophy in the computing industry. The two dozen or so "Gold Chip" customers are really the fashion setters for thousands of smaller users throughout the industry. People from these installations are the ones that give the majority of the papers at technical societies and are certainly the ones who are the most vocal when it comes to leading crusades toward new concepts and uses of computers. My point is that the two dozen or so Gold Chip customers may not represent a particularly profitable market but they represent an extremely influential market. As they go, so go a great many of the less sophisticated but more profitable customers.

The psychological factor of a company wanting to be with a winner in looking ahead and identifying themselves with future products which they may grow to is still a very important phenomena. Many small customers used to refer to STRETCH in their future plans even though they, as installations, would never grow to a size that would need such a machine.

IBM, by not having a STRETCH-like pace setter in the market, has lost a great deal of this psychological advantage in the computing industry. This has largely been taken over by the CDC 6600. The CDC 3600 is probably going to be much more of a profit maker for CDC than the 6600 will be, yet they apparently are willing to undertake this 6600 largely for the reason that it gives them the prestige of being the pace setter in the computing industry, and it will breed customers for their smaller machines.

#### VI The Legacies of STRETCH

There is still a great tendency in the Corporation to regard the STRETCH program as nothing but a large financial fiasco. A great deal of the loss attributed to STRETCH is due to two items: (1) The bookkeeping by which all sorts of development costs were charged to STRETCH alone which were later used by many other programs; and (2) by the depressing of the STRETCH price after it was initially set too high during the negotiations with Livermore. This action alone resulted in a net loss of something approaching a million dollars per machine, which certainly made quite a difference in the final profit and loss picture.

Another factor of less importance was the decision not to follow on with an improved STRETCH program. This effort could have been done fairly naturally two or three years ago, but the improved STRETCH, then called the 7034, was allowed to wither. It could have recovered some of the initial investment in overhead costs; programming, training, etc.

The truth is that STRETCH, even though it lost money as a project, was still a tremendously profitable investment for the Company in other areas. For five years STRETCH was the pace-setting machine in the computing industry. Every meeting, and practically every issue of the Trade magazines in this period, referred to STRETCH and STRETCH-class computers.

The side benefits of STRETCH such as the circuits and frames, etc., which went directly to 7090 and 7080 certainly are very well known. The stabilizing effort which this large program had upon the engineering development within the Company should certainly not be overlooked. A great deal of the effort which went into solid state circuits, etc., would have suffered even more excursions than it did had it not been for the stabilizing effect of the two important government contracts associated with the STRETCH program. (Incidentally, people tend to forget that the SSEC and the 701 were both programs of exactly this type in their day.)

Another side benefit was in the field of personnel. There were approximately 50 new people hired directly into the STRETCH program and most of them were top-notch engineers, mathematicians, programmers, etc. The Company could probably not have hired many of them without this real prestige machine which gripped the imaginations of these young men.

Another somewhat negative argument, but nevertheless a very telling one, is that if IBM had not taken the STRETCH contracts, we can rest assured that some other company would have been glad to take government money. It would, thereby, have built up its competence (using our own tax dollars) which would have long since been brought to bear against us in our own commercial strongholds.

## II Does the AEC Drive Too Hard a Bargain?

Really part of the STRETCH legacy is the feeling within IBM that the AEC drives too hard a bargain in its contracts. There is certainly some foundation to this feeling, but it is doubtful whether they are really much worse than the other government agencies.

Assuming that a new advanced computer contract were to be negotiated, it should deliberately be written as a cost-sharing venture. The inevitable side benefits in technology, personnel, and purpose would more than recover any direct loss.

VIII What About the Test Ban Treaty?

One argument I have heard against an AEC contract says, in effect: "With the Test Ban Treaty they will soon be out of business anyway." Actually, the Test Ban Treaty, assuming it is ratified, should actually increase the computing load on the AEC installations. To make comparable progress under the Treaty restrictions, much more theoretical work will have to be done to gain the results which could be determined experimentally by direct testing. The underground testing, which will continue in any event, usually calls for more computation and analysis of experiments than do the relative simple open-air tests.

There is, of course, the possibility that if a test ban moratorium continues for an extended period, that the general level of activity in the AEC Weapons Laboratories will gradually be reduced. They lived through the so-called "voluntary moratorium" which lasted almost three years in the past with very little sign of this deterioration occurring.

Another factor which should not be overlooked is the fact that both the AEC Weapons Laboratories have many other project activities going on in addition to weapons development. Many of these, particularly Sherwood (controlled Thermonuclear power) and the Rover (nuclear rocket propulsion) will require a great deal of numerical computation. So far, the use of computation design techniques is just getting a good start in these fields.

It would be extremely foolhardy to write off the AEC as being no longer worth having as a customer simply because of some political possibilities regarding a test ban at this particular time.

IX Will We Lose All Our Patents?

Another specious objection to doing business with the AEC arises with the contention that they are extremely strict in their patent position. It certainly is not the intention of the Government to prevent itself from having the latest in computing equipment simply because of some patent or legal technicalities. Certainly, we did not lose our right to exploit the 7090 because of the STRETCH and HARVEST developments done under government contract. The Patent arguments are one of those harrassing organization-man-type arguments which are impossible to answer directly ahead of time. In any particular case of interest, they can always be solved if the parties concerned want to solve them.

X Won't the AEC People Leak Our Secrets to Competitors?

Another objection I have heard is that the AEC people will take our internal Company secrets and leak them to competition. This, of course, is strictly a personal function of the individuals concerned. There was one case where this presumably happened.

If so, it could be directly traced to individuals who certainly were anti-IBM in their sentiments all along.

Actually, there is far more leakage of our intentions and our technology in ordinary technical meetings than there ever is from the AEC Laboratories. After all, AEC people are used to handling secrets of all kinds from many sources and keeping them properly compartmented. I can truthfully say that in all my years of interacting with the people at LRL and LASL, I have never received any proprietary information concerning one of our competitors, either voluntarily or unintentionally. Most of the information that one hears in AEC discussions turns out to be already widely disseminated by announcements at technical meetings, employment ads, "Electronic News" releases, etc.

#### XI Won't the AEC Scientists Insist on Designing Their Own Machines?

One of the points of view discussed at the AEC internal computing conference last year was the possibility of the AEC Laboratories buying computer components, e. g. , memory boxes, disk files, CPU frames, from different computer manufacturers and having one contractor hook them together. The proponents of this plan claimed much lower cost and better computer organization for AEC problems . Programming presumably would be done by Laboratory personnel or contracted for



separately. This point of view was not accepted by either of the weapons laboratories, although there was strong support from some quarters.

The feeling at Los Alamos now seems to be against any formal "cooperative design effort" a la STRETCH contract. I have often heard the joint STRETCH planning group criticized as being too large, too irresponsible, and its decision-making machinery too vague. Present thinking seems to favor a small, tight circle of IBM designers who are really in control of the decisions, with only informal advice from the customer's representative.

## XII In Conclusion

To save the AEC customers and hold the other "Gold Chip" customers, IBM should at once deliberately seek out a contract with the AEC (preferably with Los Alamos) for a very high speed computer having at least the performance of STRETCH on AEC problems. It should use Impact circuits, the half microsecond memory, and the large back up store. Preference should be toward a relatively simplified computer organization. "Return to simplicity" should be the motto. NPL compatibility probably will not be a major advantage in the eyes of the customer.

The contract should be gone into with the following points in mind:

1. The computer to be built should deliberately be a state-of-the-art stretcher.
2. It should be competitive and designed to get our machines back into our customers' future plans.
3. It should be a deliberate prestige-gainer. That is, a project to get IBM back into the position of being the real leader of the computing community.
4. Profit should not be the sole motive and we should be prepared to enter into a shared-cost development for the benefit of the Government. The understanding would be that the indirect benefits to be gained will really make it very profitable in the long run.

In view of the recent concern as to how to invest our capital, the above is a suggestion which will use up a goodly amount of capital, provide a purpose in life for many engineers, nail down a whole category of influential customers, push back the frontiers of computer technology and aid our national defense effort.