

*Analysis*

*STRETCH - 1961*

SAN JOSE RESEARCH LABORATORY

May 8, 1961

Personal and Confidential

Memorandum To: Dr. E. R. Piore

Subject: STRETCH Investigation

The week of April 17, 1961, when I was visiting at the new Research Laboratory, there came a call from Poughkeepsie urging me to come and talk to Messrs. C. Adams and R. Meagher concerning STRETCH. As I get the story, they were commissioned by IBM CHQ to make an independent investigation of STRETCH and come up with an evaluation of the machine and the project which created it.

I was somewhat disturbed by this investigation and the implications which one could draw from it, especially were one to give "investigation" its usual newspaper definition. Secondly I was concerned that outsiders would have a hard time extracting the truth at this late date, particularly in view of the fact that so few of the original crew are still on the STRETCH project. I discussed my concern with Dr. Gil King last week. He suggested that I write this memo to you and assured me that it would be received with proper attention and confidence.

As you know, my relationship to the STRETCH project has been somewhat unique. I was an Associate Group Leader of the Los Alamos, Theoretical Division, when the proposal was first made in 1955. I served on the LASL evaluation board which gave the original contract to IBM and was a member of the joint IBM-LASL Mathematical Planning Group at the time I left Los Alamos to join IBM the summer of 1957. At Poughkeepsie I was a member of the

Product Planning Group attached to STRETCH until I left to join FSD in June 1959. For the last two months of the period I was in charge of what was left of the Product Planning Group. I have attempted to set down here, in non-technical narrative style, my impressions of what happened and where the main problems lay.

Before making any comments which might be construed as negative, I want to say that in no sense should IBM consider the STRETCH project as a mistake for the Company. By it IBM has maintained the unquestioned leadership in the eyes of the computer community for over five years. It has created a new word in the language. Terms like "STRETCH class", "pre-STRETCH", and "post-STRETCH" are used freely in the literature to describe the upper end of the computer spectrum.

I need scarcely mention the tremendous side benefits derived from the use of STRETCH circuits, memory and logic in the 7090 and other programs.

Even in its defeats (e. g., failing to meet the original delivery schedule) the Company has conducted itself with dignity and did not resort to such tactics as shipping an incomplete or untested machine.

Certainly an evaluation is called for now that the machine is shipped and the program can be examined with more detachment. In making judgments, however, it is important to keep in mind the time scale of events. Many of the key decisions were made long before there was any experience with transistorized machines, and long before the existence of a working 2-microsecond memory. Also recall that external competition and IBM's internal organization has turned over several times during the life of STRETCH.

The STRETCH program was beset with problems in the areas of: (1) Technology, (2) System planning, (3) Financing, (4) Market or product planning, (5) Internal organization, (6) External pressures, (7) Contractual negotiations, etc. Many decisions were made in each of these problem areas and a number of mistakes were made. One can expect and forgive an occasional lapse in judgment.

However, when I look at the cause of such individual problems, I find that there seems to be a recurring theme behind every one of them--the failure to consider STRETCH as a complete systems design in any consistent sense.

(By "systems design", I mean considering the computer as an organic whole, in which a change in one section must be considered for its change upon the entire performance of the machine. The logical design, the hardware technology, the input-output devices, etc., must be considered in their entirety as applied to the problems which are to be solved. Inherent in systems design is an iterative procedure: goals, models, experimentation, evaluation, modification, new models, new experiments, etc. The anti-thesis of systems design is the "black box" approach in which a set of individual devices are built to fixed specifications and independently hooked together after the fact, with no iteration nor chance to modify the design after evaluation.)

In reviewing the project it becomes clear to me that there was a lack of consistent and agreed-upon goals on the part of the many people working on the project in its early formative years. It seemed as though everyone had a different goal for STRETCH to achieve. The AEC proposal and contract, which should have been the rallying point, was often completely ignored. To a casual observer or even someone in the project at the time, it was difficult to tell what these conflicting goals were.

Another recurring attitude was the unwillingness to consider sufficient detail before making decisions. There was a "happy optimism" concerning the outcome of all the more difficult points, while a great deal of time was spent on relatively simple or trivial points. Of course, one cannot criticize over-optimism in a new project. However, as the project wore on, this optimism should have been replaced by a better knowledge of the facts. This was brought forcibly to me when the first results of the Cocke-Kolsky STRETCH timing simulation program came forth and demonstrated that the machine at that time (January 1958), even with relatively optimistic assumptions, was only half the performance which people were assuming that it was.

Enough of these generalities. Let us go back and examine the growth of the program chronologically.

In early 1955 the AEC Laboratory at Livermore put out a request for bids to the various computer manufacturers to build a computer considerably better than any then in existence. IBM prepared a machine design sometimes called "the Livermore 2-Megacycle Proposal". There was considerable jockeying as to whether or not this proposal should be made. In fact, the description of this action as told to me by one of the participants sounds more like a game of cops and robbers rather than a serious proposal to a government laboratory. The net result of this half-hearted effort was that IBM lost the contract to Sperry Rand. The machine which they proposed and later built and delivered is now known as the LARC.

In the months that followed, the conviction grew that a machine superior to LARC by as much as a factor of 10 would be possible. The decision was made about September 1955 to propose such a machine to the AEC Los Alamos Laboratory. An internal IBM memorandum at the time calls for: 10-megacycle transistors, 1/2-microsecond memories, floating add times of 0.6 microseconds, floating multiply 1.2 microseconds, fast tapes providing 2,400,000 bits per second, and delivery in October 1958 (36 months).

A presentation (which I attended) was made at Los Alamos September 20, 1955, quoting a price of \$3.6 million. After serious consideration, the AEC decided that a request for bids for such a large machine would have to be submitted to more than one computer manufacturer. Such a request for bids was sent out (based upon the STRETCH specifications) in January 1956. The AEC received proposals from four computer manufacturers, including the IBM STRETCH proposal. The proposal was essentially the same except that the price figure was raised to \$4.3 million. What had happened during this period was a growing realization within IBM of the total cost of such a program. The original price of \$3.6 million was known to be far less than the project would cost. The increase in price to \$4.3 million represented some sort of promise to reduce the development costs to IBM to a "reasonable" level. (This change caused considerable embarrassment at LASL, where \$3.5 million (and no more) was already authorized.)

May 8, 1961

Note: This tendency to "tighten the belt" spasmodically without regard to the technical needs was a recurring phenomena for the next several years, and was one of the main factors influencing the final outcome of the program.

An AEC technical committee made a study of the four proposals presented to Los Alamos. Since I was a member of this study team and working for Los Alamos at the time, I cannot discuss any details of the evaluation. However, I can say that at least two of the other proposals (in hindsight) were more realistic than was IBM's.

Be that as it may, Los Alamos finally decided upon IBM, largely on the basis of faith. After a surprisingly long period of negotiations, the AEC contract was finally signed in November, 1956. (The delivery date was now May 1960--42 months.) The contract called for a joint IBM-LASL Planning Group which would meet for one year (December 1956-December 1957) to decide upon the logical design of the machine. The contract was quite inexact in spelling out procedures which this Planning Group should follow in considering and adopting features. The intention in being vague was not to tie the hands of the engineers and logical designers with preconceived notions.

Although the study group did not get underway formally until December, 1957, meetings were held throughout 1957 involving the same people which were to make up the Planning Group. (There was one important face missing--that of Gene Amdahl who had left IBM in the meantime.) Meetings of this group were held with the usual informal friendly exchange of ideas which characterized the earlier proposal discussions. Somewhere about April, 1957, however, there was detected at LASL a distinct cooling of the relationship. In March and April, a period of a month or six weeks went by with no contact at all between LASL and IBM. Actually during this time a study group known as the "Three-in-One Committee" was absorbing all the planning effort. The goals of this study were to come up with a plan for integrating into one development program three completely different computing systems: 1) The Los Alamos high speed scientific computer;

May 8, 1961

- 2) A low-speed commercially-oriented computer ("Basic"); and
- 3) The high-speed HARVEST data processing system. A report of this plan is dated June 1, 1957.

In a staff report published in September, 1957, a schedule is spelled out which states that the "Basic-Model Machine" was to be designed by March 1958 and released to Manufacturing September '58 through July '59. The same schedule shows the planning and logical development of "Sigma" (the Los Alamos machine) as taking place from December '57 to December '58, and its release to Manufacturing from October '59 through February '60. The HARVEST system's planning and logical development was to continue on through July, 1959, at which time its model construction would begin. (HARVEST was never considered as a released item.)

Notice the inference drawn from this schedule: The "Basic Machine", which was to include all of the address logic, the variable field length operations, most of the indexing operations, and the input-output controls--in other words, all the connections between the machine and the outside world and most of its logic--was to be frozen early in 1958, a good nine months before the official IBM-LASL Planning Group was due to complete its deliberations. LASL was not informed of this approach except indirectly. They did, of course, become aware that there seemed to be an increasing number of "untouchable" constraints being insisted upon by IBM in the design of STRETCH. The explanation or justification for these decisions was never quite made clear because stating the true reason, i. e., the development partially under the contract of a computer for commercial data processing, was not something which the AEC could be expected to embrace whole-heartedly.

Another factor which was overlooked even in the schedule which I have just quoted was that in truth the HARVEST planning and logical development was considerably further advanced than that of either Basic or Sigma. There had been contracts for two or three years between NSA and IBM to examine the requirements and logical structure of a computer such as HARVEST. There existed as early

May 8, 1961

as May of 1957 a large and fairly complete manual describing a machine system for HARVEST. A revised version of the manual was published in November of 1957. The decision to use the same Basic Computer for both HARVEST and Sigma placed additional constraints upon the STRETCH logical design. Again the AEC was only vaguely aware of these conflicting goals.

Another ingredient to be poured into this pretty confused pot was the fact that IBM had a reorganization of the STRETCH program at this time, in which the computer project was passed from Research to the Product Development Laboratory, but many of the original planners were in a "Product Planning Group" which reported through marketing channels or in Corporate Research. This presented a confusing picture to the Los Alamos representatives on the Joint Planning Group. Again, they were only vaguely aware of the fact that half of the people representing IBM were directly concerned with the construction of STRETCH and the other half were largely concerned with "keeping track of the other half", or squeezing a commercial development out of the program.

It is interesting at this late date to read from the original "Three-in-One Committee" report. Their assumptions, in view of our present knowledge, seem quite naive. For example (the underlining is mine):

"It is possible to set forth a workable 3-in-1 system meeting our goals, with the possible exception of price."

"The HARVEST system is readily separable into a specialized unit H and a general purpose unit B. If the system is to be built, this separation should be preserved."

"Integration into the 3-in-1 system limits the logical and variable field arithmetic speeds obtainable for the Sigma System. If faster speeds for these operations are needed, the 3-in-1 approach must be abandoned."

"An independent senior technical computer could perform most of the functions of the B & S, and offer the possibility of adding equipment to increase logical and variable field length

operation speed. Preliminary investigation indicates that such an independent machine could be built with slightly less hardware than B & S. "

My only comment is the "senior technical computer" is what the project was supposed to be building in the first place. The Three-in-One approach was indeed finally abandoned about six to nine months later--not, however, before its serial-parallel logical structure was imposed upon the entire STRETCH design.

Another political development which occurred during late 1957 was the attempt to bid a derivative of the Basic computer (called "Junior STRETCH") for the 465L Air Force Program. After a fierce interdivisional fight within IBM, the decision went in favor of bidding a Kingston design based on the RTA machine. This skirmish had no direct effect upon STRETCH except that a large percentage of the effort of the engineering planning group was spent for many weeks freezing and re-freezing the design of the Junior and Basic computers. This delayed the eventual cancelling of the Three-in-One concept.

During the fall of 1957 the Joint IBM-LASL planning group meetings became longer and more feverish. Arguments built up and tempers mounted. Most decisions were reached by the simple compromise of including both proposals in the machine. Only a few voices were raised protesting the complications which all these compromises might impose on the hardware, but they were shrugged off because there was no way of evaluating the cost of anything-- either dollars or performance. When the full horror of the engineering complications began to be felt during early 1959, as the detailed logical design of the boxes was being laid out, it was then "too late" to reconsider the logical structure of the computer--just as it had been "too late" in 1958 and 1957. In my opinion, this unwillingness to ever take a second look at the over-all systems design was one of the most serious weaknesses in the STRETCH program. Because of it, effective use was not made of the official planning period in 1957 nor later. Preconceived decisions based on other commitments (or other imagined commitments) always prevented a free consideration of all aspects of the design.

Most of the detailed complications in the STRETCH logical design arise from the inclusion within a single machine of two diametrically opposed philosophies of operation: One is the principle of parallelism, i. e., having many components of the system running at the same time to gain high performance. The other is the principle of interruption, in which any unusual event can trigger a change in the instruction stream at any time. The former results in a logical design such as the "look-ahead system" of STRETCH in which as many as seven instructions can be in different stages of execution at once within the computer. The interrupt system, on the other hand, states that an interruption can occur at any time to terminate the stream of instructions and restart from a new location. The results of a given interruption must be completely independent of the particular sequence of orders which is being executed at the time. In STRETCH an interruption must stop the flow of instructions, allow those to be completed which are legitimate, and cancel out the effects of those which should not have been started. Of course, one can say that either of these philosophies is good--they are both highly desirable features within a computer and must be included in some form or other. The problems arose in STRETCH when the logical implications of the particular combination chosen were not worked out soon enough to allow an examination of their real cost. I like to visualize the interaction of these two opposing philosophies as the clash between a large warm air mass and cold air mass in the atmosphere which results in a line of thunderstorms building up to tremendous heights. This is exactly the kind of effect which one sees in the I-box of the STRETCH computer. The control and logical complications required to recover precisely from every type of interruption under any condition of look-ahead overlap is really a tremendous accomplishment for the engineers who laid it out.

More generally, a serious mistake was made when the senior planners who decided upon this particular machine philosophy were not forced to draw out the circuit diagrams down to a level of, say 15 or 20 blueprints. They instead stopped at the level of about three blueprints and left the further breaking down of the design to junior engineers. The latter saw the horrible complications months later, but were not capable of going back and re-examining the original assumptions. If the senior planners had been forced to lay out all the

data paths and all the main control lines, they would have seen the complications which arose soon enough to be able to do something about them. This principle of follow-through or iteration of design is an extremely important one in systems engineering which was largely missed in STRETCH.

The above battles on the logical and political scene were more obvious but not necessarily any more serious than another simultaneous series of battles which were taking place on the financial front. It was realized from the beginning that STRETCH would cost considerably more than the AEC contract would bring in. However, the full realization of this did not become apparent until the engineering build-up during 1957 was taking place. The result came in the form of a series of financial crises concerning the budget. In desperation the STRETCH managers sought means to stay within the budget while still doing the technical job. One way chosen was to split off a number of the peripheral functions of the project such as the basic circuits group, human engineering, programming, etc. Splitting off the circuits group was particularly disastrous. In a short time the new separate group was imposing standards on STRETCH (e. g., the single card) which would have completely ham-strung the machine. The result, again born of desperation, was the formation of another rival circuits group within the STRETCH project. Eventually this led to the development of the so-called "double card" which was designed around operational elements rather than individual "and/or" logical blocks. The double card was essential to achieve the short lead lengths and packing densities required to build a machine as large as STRETCH. This "war" between the circuit groups had far-reaching effects upon the whole machine design and certainly drove the costs way up.

Another area which suffered during these financial crises was the one which always suffers on computers, namely, the input-output devices. Pressed for funds, the STRETCH project had relatively weak leverage in trying to obtain input-output device developments from other laboratories and Divisions. The STRETCH disk file is a good example of such a starved item.

Another: The STRETCH console program was almost killed at one point. Systems planning was certainly difficult because of the dispersal of the development of each of the major components. As late as May of 1959 the planners in Poughkeepsie had only a vague feeling as to the status and actual operation of the chain printer or STRETCH disk which were two of the large components of the operational system.

On several occasions during '57 and '58 orders came from White Plains: "Cut down the transistor count"; "Get the cost below such and such". Of course there were always those who wanted to cancel the project altogether because of its high cost. I have saved several memos which were written at one time or another during these financial clashes. One of them dated June 1958 gives a check list for "simplifying" Sigma. Through the 16-page memo practically every single desirable feature of the machine is suggested for removal one by one. Along with each feature is an estimate as to the number of transistors which would thus be saved. Interestingly enough some of the features which were actually removed at that time are now being put back as part of the machine improvement program.

If I were to summarize the most serious mistakes or sources of difficulty I would list the following:

1. The loss of Gene Amdahl at the critical early stage of the project.
2. The premature removal of the project from Research, placing it in the Product Development Laboratories, which resulted in the "circuit war" mentioned above. Had the program remained in Research or at least been kept intact, questions such as circuit standards could have been resolved without the heavy emphasis upon finances and the requirements of other programs.

3. The wandering away from the true original purpose of the STRETCH program, which was to develop a senior scientific computer of the next generation. The substitution of other (often obscure) goals by key personnel was a major factor.

4. The "Three-in-One" concept which, although it was later cancelled, left a legacy of the parallel arithmetic unit and a serial arithmetic unit, full word and half word instructions, etc. I directly attribute to it such strange inconsistencies as the fact that STRETCH now can perform a floating multiplication three or four times faster than it can branch on the result of an arithmetic operation. (Incidentally, this same philosophy is again being applied to the 8000 series--but that is another story.)

5. The unwillingness on the part of the Corporation to commit its funds to the project early in the game to solve the basic technical problems while there was still time to solve them. Far more money was spent on overtime and crash efforts in 1960 trying to catch up after it was too late than would have been needed had the money been spent gradually over the period of three years. Actually, the Three-in-One concept itself was brought about as the result of financial pressure, that is, the requirement to obtain a "wide market basis" for the project to stave off managerial action.

6. The deliberate splitting off of people within the Company who were familiar with all problems and their requirements and putting them in Product Planning or on other projects, e. g., HARVEST or Lightning rather than allowing them to work directly on the AEC machine.

7. An iteration upon the total systems design was not made when it was realized that some of the basic assumptions were leading to difficulties.

8. A sufficient systems evaluation study or simulation of operation was not carried out. The work of this sort which was carried out was too late and generally disregarded.

9. The Joint AEC-IBM Mathematical Planning Group was not organized properly. Proper decision procedures were not spelled out, and the meetings frequently degenerated into free-for-alls with no one clearly designated to make final decisions. This was coupled with the fact that it operated in a vacuum much of the time concerning the nature of the logic and circuit constraints being placed upon the system by other commitments.

10. The lack of early programming support for STRETCH was quite serious. The only sizable group of programmers assigned to STRETCH for a long time were really the HARVEST "Farm-Boy" group. When Applied Programming finally came into the picture (they did not do so until early 1959), there resulted a long series of battles in which they objected to everything (perhaps justifiably) and set about changing the names of all the instructions and operating philosophy behind the computer system. These differences between the groups were worked out in time to a quite satisfactory conclusion, but notice this was done in 1959, not 1958 or 1957.

You will notice that I avoided (except in one case) mentioning personalities in the above narrative. Needless to say any such large project is strongly influenced by the personalities on it, above it, and outside it. I will deviate from my rule again and mention one other name. I hope that as the investigations proceed, there will not be the tendency to blame everything bad in STRETCH on Steve Dunwell. In my opinion he did a tremendous job in an impossible situation. If he is to be blamed for all the bad things, he should also be credited with all the good, which certainly outnumber them manyfold.

Dr. E. R. Piore

-14-

May 8, 1961

Again I want to say that these represent my own observations from a limited perspective. I would be glad to discuss them further if it seems worthwhile.

H. G. Kolsky

HGK:lc

cc: A. G. Anderson  
G. W. King