STRETCH by S G Campbell

POUGHKEEPSIE
Department 4 /5

April 12, 1960

Memorandum to: Mr. C. A. Merritt

Subject: Lectures at U. C. L. A., May 23-25, 1960

The enclosed correspondence relating to subject lectures is largely self-explanatory. The brochure has been widely distributed by U. C. L. A. As indicated in the schedule, the IBM machines are given three days (six hours per day, except for the subtraction of 45 minutes the first day), followed by the Sperry Rand LARC, Ferranti ATLAS, Bull Gamma 60, and University of Illinois Computer. It is noteworthy that my presentation on Stretch and Harvest will be followed by the LARC presentation given by Dr. Sidney Fernbach of the University of California Radiation Laboratory at Livermore, California. It should be regarded as fortunate that this presentation is in the hands of a serious and responsible scientist, rather than a Sperry Rand salesman. Dr. Fernbach will give a careful and balanced appraisal of LARC.

It is probable that the audience will be composed of students, company representatives, and perhaps a few (but very few) representatives of Government Installations. An examination will be given at the close of the course, and students will be graded. The company representatives will largely represent users of large scale computing systems, but undoubtedly our competitors will be represented thoroughly. The lectures will be purely technical, designed for experienced professionals, and will approach giant computers from the point of view of the user.

Normally, the above conditions would not be pertinent to the clearance of a set of lectures to an unrestricted audience. In this case, however, there are a number of unusual considerations:

1. I am scheduled for 17 hours 15 minutes of lecture and discussion. A script for this performance, single spaced on an IBM typewriter, would require more than 400 pages at my lecturing pace, even assuming that I could or would lecture from a script. Preparation of such a detailed script is impossible, and its use would leave a disastrous impression even if it were possible.

2. Each of the other lecturers is associated directly with a University; I am the only one directly associated with a Manufacturer. The others will give good technical expositions which are detailed, fair, and well balanced. I consider it important to IBM that we be represented by lectures of comparable quality, and that we not take advantage of our special position.

3. In many cases, principles of good exposition will be the main consideration. For example, how much I say about the Stretch Look-Ahead probably depends more upon precepts of good exposition than upon Company Confidential restrictions. Published information is probably adequate (but is also partly incorrect).

4. Marketing Considerations. I have been discussing Marketing Considerations with Mr. Dura Sweeney in Poughkeepsie and Dr. Lou Robinson in White Plains, and will remain in close contact with Marketing through the presentation. We have made arrangements for me to work with Marketing representatives during the presentation to ensure that our marketing interests are protected. Arrangements have also been made to handle any matters involving press coverage. By this, I mean that care will be taken to protect our position in terms of a Marketing situation which may change right up to the time of the lecture. This does not mean that I will deliver a sales pitch.

5. Company Confidential Information. The following list of cleared publications covers the entire Stretch System and much of Harvest. A great deal of the material I will cover is naturally included in these papers. In some cases, considerably more detail is included in the paper than I will want to present. However, the material presented in these papers is inadequate for my purposes for several reasons:

   a. With the exception of 2, 11, and 14, the papers are concerned with specific sections of Stretch and Harvest. There exists no unified approach to the total system, much less to the Operating System, which is the standpoint from which I will approach these machines. Hence these papers do not add up to an operating Stretch System. Most of these papers have been cleared because the writer made the effort to obtain clearance for a specific paper dealing with a specific area.

b. Many of the papers represent objectives rather than accomplishments. I will be speaking about a system I will have actually operated (my maintenance programming group has been operating Stretch for three weeks already). If I am required to stick to published data, just because it has been published, at a time when we have actually tested the machine, then our position is bound to be foolish.

c. The papers in question must be presumed known to the bulk of the audience. I propose, therefore, to commence by stating overall design objectives in terms of these original papers, and relating this to the machine structure as it has actually evolved. The emphasis here will be on overall machine performance, rather than performance for specific operations, and design objectives will be treated more as objectives for the total operating system than as, say, objectives for a specific multiply speed.

d. This published material can also serve a valuable purpose as reference information. Whenever possible, I plan to refer questioners to the published literature for further detail on specific areas of the machine. For example, I do not regard a detailed question on circuitry as pertinent to my subject matter, and would refer such questions to Erich Bloch's paper (12).

1. E. F. Codd, E. S. Lowry, E. McDonough, C. A. Scalzi, "Multiprogramming Stretch: Feasibility Considerations", Communications of the ACM (November, 1959).

2. S. W. Dunwell, "Design Objectives for the IBM Stretch Computer", Proceedings of EJCC (December, 1956).

3. F. P. Brooks, Jr., "A Program-Controlled Program Interruption System", Proceedings of EJCC (December, 1956).

4. F. P. Brooks, Jr., "The Execute Operations - A Fourth Mode of Instruction Sequencing", Communications of the ACM (March, 1960).

5. W. Buchholz, "Fingers or Fists?", Communications of the ACM (December, 1959).

6. F. P. Brooks, Jr., G. A. Blaauw, W. Buchholz, "Processing Data in Bits and Pieces", IRE Transactions on Electronic Computers, EC-8, No. 2 (June, 1959); (Also IBM TR00.01000.674).

7. W. Buchholz, "The Selection of an Instruction Language", Proceedings of the WJCC (May, 1958).

8. G. A. Blaauw, "Data Handling by Control Word Techniques", Proceedings of the EJCC (December, 1958).

9. G. A. Blaauw, "Indexing and Control Word Techniques", IBM Journal of Research and Development (July, 1959).

10. H. K. Wild, "The Input-Output Devices of the Stretch Computer", paper presented at Auto-Math '59 Exhibition, Paris (June, 1959).

11. W. Buchholz, "Design Objectives for the IBM Stretch Computer", New Computers: A Report from the Manufacturers, ACM (March, 1957).

12. E. Bloch, "The Engineering Design of the Stretch Computer", Proceedings of the EJCC (December, 1959).

13. J. Cocke, H. G. Kolsky, "The Virtual Memory in the STRETCH Computer", Proceedings of the EJCC (December, 1959).

14. P. S. Herwitz, J. H. Pomerene, "The Harvest System", to be presented at WJCC (May, 1960).

Papers which have to the best of my knowledge not been cleared for publication and which contain information I wish to use include:

a. STRETCH Data Processing System, Preliminary Manual of Operation (June 1, 1959), IBM Confidential.

b. HARVEST Data Processing System, Preliminary Information Manual (December 1, 1959), IBM Confidential.

c. W. C. Stetler, SIGMA Computer Memo #26A, "Look-Ahead System", (January 19, 1960). No classification indicated.

d. R. T. Blosk, SIGMA Computer Memo #36, "Stretch Computer Instruction Unit: Design and Performance Goals", (March 9, 1960), IBM Confidential.

e. H. C. Montgomery, "Floating Point Division in SIGMA - The Algorithm Used", File Memo, 7030 Engineering Planning, (June 9, 1959), IBM Confidential.

f. H. C. Montgomery, "Data Flow in Floating Point DIVIDE", File Memo, 7030 Engineering Planning (June 30, 1959), IBM Confidential.

g. H. C. Montgomery, "Floating Point Multiply in SIGMA", File Memo, 7030 Engineering Planning (June 15, 1959), IBM Confidential.

h. R. J. Bahnsen, "Description of Checking in the Sigma System", SIGMA Computer Memo #28, (February 9, 1959), IBM Confidential.

i. W. R. Stringfellow, "Maintenance Scanner", File Memo, 7030 Maintenance Reliability, (September 29, 1959), IBM Confidential.

j. STRETCH Data Processing School #4218, IBM Department of Education, (February 16, 1959).

k. Preliminary Stretch Programming Manual, IBM Applied Programming, (February 1, 1959), IBM Confidential.

l. R. P. Fletcher, G. E. Russell, F. C. Tung, H. K. Wild, "Basic Exchange Machine Specifications", Exchange Memo #32, (December 4, 1957), IBM Confidential.

m. L. R. Johnson, "A Description of Stretch", (December 10, 1959), IBM Confidential.

Of these, (e), (f), (g) are probably covered, for my purposes, in (12) preceding; (j), (k) are not really pertinent - have been superseded by later developments (and never did have much pertinence); (l) is covered by (10) preceding; (h) can be covered through the original Hamming paper; and (i) is desirable information, but could be eliminated with the remark found in Kolsky (13) on this subject. However, (d) contains important information (although most of the information in (d) is not required). Obviously, (a) and (b) are particularly important, since they contain the instruction sets and formats which are required for any sensible presentation of the Stretch and Harvest Systems; (m) also contains important material, although some of it is incorrect. I have been informed by Marketing that, depending upon decisions of when and how to market Stretch, they will probably make available the manual in its entirety before my lectures. Publication of the

Stretch manual is therefore primarily a marketing consideration, and I have made arrangements to work closely with Marketing on all facets of the presentation. If arrangements have been or will be made by Marketing to clear the manual, however, this will automatically clear a great deal of the information I require, and should simplify your task as well as mine.

With respect to the Harvest manual, there is no particular reason for clearing it, since there are no present plans for more than one machine. Harvest must therefore be presented as a "one shot" computer, and I should be personally opposed to clearing the entire Harvest manual at this time. What I do require from it is the description of the Harvest units and of the Harvest instructions, material which has been implicitly cleared in the Hervitz-Pomerene paper (14) above. The Harvest customer must also be considered in the Harvest presentation, and I will go over the Harvest presentation in detail with representatives of the customer to make certain that customer interests are protected. Although the situation with respect to Stretch is different, I shall also consult with Los Alamos as a matter of good customer relations.

While the computer itself is covered rather well by previous publications, the same is not true of the programming system. This is largely because the programming system comes later in time than the computer, and because we have held back any such publication pending announcement of Stretch - a detailed description of the programming system (the "programming package") is considered by Marketing to be an essential part of Stretch announcement plans, and a rough draft of such a write-up has been given to Marketing for this purpose. Since programs are not patentable, and are freely copied throughout the Industry, and since programming ideas are pretty well held in common throughout the more advanced programming circles in this country, the only competitive value to programs is the programs themselves - they can be easily copied. The information on the programming system which has been presented thus far is primarily included in the announcement made to SHARE last year (describing STRAP, the Stretch Assembly Program, in detail, and outlining plans for a Macro-Language, a Master Control Program, and a Fortran program for Stretch, as well as the "704 package" designed to make Stretch more useful to customers with 704's installed). STRAP-I was in fact written by Los Alamos, and we have written permission from them to use this material so long as credit is properly assigned. If Stretch is announced before my lecture, the announcement

will include further detail on the programming system, as indicated above. In any event, I will not discuss the details of how the programming system works, but rather its use and the facility it gives the user. I believe that the programming system is a very powerful marketing tool, and plan accordingly to coordinate the discussion of it very closely with Marketing. There is no question but that the manner of discussion will depend strongly upon whether Stretch has been announced or not. In the case of Harvest, the programming system is more customer-oriented, and we must respect customer preferences regarding release of information. This does not seem to pose any problem; the programming system is unclassified so far as government security is concerned.

6. General Scheme of Presentation. Using the published papers as an implicit starting point for my lectures, I expect to begin by bringing these early systems descriptions up-to-date with a description first of data flow, then of major systems components, and, finally, of the detailed instruction set and formats. This will occupy the first day's lecture, particularly since the instruction set and the instruction and data formats should be considered in some detail. The main difficulty would appear to be the parallel disk file. This situation is covered by my enclosed letter of April 5, 1960 to Mr. A. F. Shugart. My own feeling in this matter (and in general) is that the decision as to what to say about a specific device (such as the disk file) should be based on the latest possible information. The alternative is to use Contract Specifications, which have mostly been published, but which may not reflect the real situation. This could be foolish only a few days before scheduled delivery.

The second day will be devoted to a description of the programming system - STRAP-I, STRAP-II, Master Control Program, Macro-generator, simulation and utility programs; timing programs. This will be followed by simple programming examples to illustrate actual characteristics of the system. No clearance difficulty is anticipated here.

The third day will be devoted to the Harvest computer. To remove a common misconception, neither Harvest hardware nor the Harvest Programming System has any government classification; only the problems are classified. A rather detailed systems description has been cleared for Harvest already in the paper scheduled for presentation May 3, 1960 at the WJCC by P. S. Herwitz and J. H. Pomerene,

both of IBM. As in the case of Stretch, I propose to describe the Harvest data paths, the Harvest units, then the Harvest instruction set and formats. This will be followed by a description of the Harvest Programming System and an example of a non-classified Harvest problem. (The Herwitz-Pomerene paper contains such an example.)

7. It is particularly important to observe that the lectures will commence only a week before the Stretch machine is scheduled for delivery to Los Alamos. By this time, I expect that we will have had extensive operating experience with the System, will have run actual problems, will know a great deal that we do not now know. At this point, it would be foolish to give a Systems description which does not match the known operating characteristics of the actual system, despite the fact that such a description had been cleared for presentation. This does not involve many points of concern, but it does involve some; certainly, some of the published figures do not appear realistic. In the case of Marketing considerations, we have arranged to take into account the situation as it exists on the day of the lecture. I feel that we need a similar arrangement for such questions as actual machine speed, disk capacity and performance, tractor tape performance, and Harvest fast memory performance. The advantage of describing an existing, operating system in realistic terms can offset performance which may in some cases be somewhat less than indicated in previously published information. It will be impossible to give a true feeling for a system actually in operation if I am required to use fictitious figures to describe it. My emphasis, however, will be on the total capacity of the system, not the details.

8. There will be a campaign to dramatize the Stretch System as a powerful IBM Engineering Achievement. This seems a valid and important function, but I do not feel that it relates to my assignment of presenting a serious and difficult technical lecture at a major University. I know and have talked with many of our larger customers; all of them have very capable, tough, technical people, who will not be impressed by the publicity or by the test problems. I feel that my role is to talk to these people - to the senior technical personnel, who are least affected by the usual publicity.

I am preparing a detailed outline of the presentation I plan to make and will send this to you as soon as possible. The purpose of this memo is to give you a general outline of the situation, and of the intended lectures. Information would be appreciated as to those areas which are considered most likely to raise clearance problems. With regard to publication of the

material developed in the course, I assume that this would require approval, regardless of prior clearance, but would like your opinion as to how best to proceed on the matter. (See enclosed letter of February 29, 1960, from C. B. Tompkins.)

Since activities in my own department require full time, I am necessarily developing the lecture material evenings and weekends. However, I will get the information you require to you as quickly as possible.

S. G. Campbell
Advanced Systems Programming
Manager

SGC/img
Encl:
cc: Dr. C. R. DeCarlo
Mr. S. W. Dunwell
Mr. H. T. Marcy
Mr. J. C. McPherson
Mr. R. E. Merwin