REPORT ON EUROPEAN TRIP

A. S. Householder
Oak Ridge National Laboratory

This report may well appear disjointed. In the course of nearly five weeks spent in visiting many different establishments and talking with many different people, many suggestions and odd bits of information were picked up and it will take some time to assimilate, evaluate, and develop them. A few general observations may be made. Beyond that, I can only mention some of the things seen and topics discussed at each place in turn.

To begin with it may be of some interest to indicate how it seems, in the light of present hindsight, the itinerary should have been organized. Across the width of the Atlantic it is difficult to determine what places are most worth visiting, or for how long. Actually these variables were not entirely subject to control. Thus Wheeler at Cambridge, Fox and Goodwin at National Physical Laboratory, were not available before the last week in September, and a visit to Rome had to be canceled at almost the last minute. This left me with more time at Zurich than was needed and even then I missed seeing Stiefel who, unexpectedly, was away all that week.

The most regrettable omission was a visit to Stockholm to see the Besk. This is the only machine now operating abroad whose speed is comparable to that of the Oracle, with division at 600 μs. Both Darmstadt and Basel are on the way from Amsterdam to Zurich. While neither has a digital computer, both have very active and able numerical analysts, under Alwin Walther at Darmstadt, and A. Ostrowski at Basel. Also Darmstadt has a considerable variety of analogue equipment.

Neither Rome nor Paris is perhaps of great interest, computationally, at the present moment, although the National Institute for the Application of Calculus at Rome is one of the older and more important centers for numerical analysis on the continent. The Italian government has purchased a Ferranti computer which is now operating at the Ferranti plant in Manchester, and is to be shipped to Rome this fall. In Paris, Couffignal is to get an Elliott 402 this fall, and at the Laboratoire Centrale de l'Armement a small experimental digital computer is in operation and another is being built. Had the trip been made later both places would have been of interest.

I spent only one afternoon at Manchester, all of it at the University. I learned too late of two things, one being that Ferranti is developing a low-cost (£30,000) machine using nickel delay lines. The other is that computing activity is developing at the College of Technology, presumably encouraged by the new chancellor, Bowden, whose book Faster Than Thought is one of the best books now available on automatic computation.

Elsewhere in England, probably the most interesting place I did not get to is the Armament Research Establishment, Fort Halstead, where a Ferranti computer has been in operation since early this year. The atomic weapons establishment is also to install a Ferranti, but at a somewhat later date.

Computing, and computing machine development, in England and in Europe, is largely governed by an important economic consideration, that human resources are abundant and cheap, whereas money is scarce and equipment expensive. In Holland, for example, a job would pay roughly N guilders if here, in this country, it would pay N dollars. A secondary factor, which may be somehow related to the first, is that such activity seems to be relatively more common in educational institutions and carried on for pure research. Thus, both the construction, and the subsequent use, of Edsac I at Cambridge, were primarily for research purposes, and although the Lyons Tea Company contributed some funds for the construction, they apparently did not exert any control. At the University of Manchester, a transistor machine is being constructed, purely for purposes of engineering research on the uses of transistors. This seems to contrast with the situations at Harvard and at MIT, for example, where the several machines were built largely with military funds for specific purposes. The only comparable case in this country seems to be the University of Illinois where the Illiac is used exclusively for research purposes by the University.

So far as I could learn, there is no commercial manufacturer of computing machines anywhere on the Continent, and in England there are only three. These are Ferranti, Elliott, and English Electric, and English Electric seems to be half-hearted, at best. The Ferranti machine, of which 6 have been completed, 5 delivered and one soon to be delivered, is basically an engineered version of the Manchester machine. They are also developing, as mentioned above, a new type of machine at lower cost. Elliott Brothers are building a machine, the 402, which sells at £25,000, one being in operation in their plant at Borehamwood, and three in assembly. They were the first to develop the nickel delay line, which is cheaper and sturdier than the mercury delay line, and which Ferranti has now taken up. They have further developed three types of plug-in units from which they say they can construct a machine to almost any desired specifications with accurate costing from the drawing board. English Electric is making the

Deuce, an engineered version of the Ace at National Physical Laboratory. One of these is to go to NPL; one goes to the Royal Aircraft Establishment; one or two others are being built but with no definite purchasers in sight.

Whatever may be the reasons for it, the machines now operating, and those under development, are generally relatively slow. Only three of those under development promise anything like Oracle speeds. These are Edsac II at Cambridge, a machine under development at the Mathematical Center at Amsterdam, and a machine now being completed at the University of Manchester. Nevertheless, by use of ingenious logic, components that are intrinsically slow are being made to yield the utmost in speed and flexibility. Moreover, I get the impression of many diverse lines of development going on more or less independently. Elliott and Ferranti are exploiting the nickel delay line as a main store. Harwell and Manchester are building transistor machines, with magnetic drums for the main store, but with markedly different logics. Amsterdam and Cambridge are building machines with magnetic cores for storage, but Cambridge is using cores also for switching elements. Cambridge is further utilizing a logical device due to Wilkes, of constructing machine orders out of a limited number of "micro-orders", which provides them, they feel, with almost infinite flexibility in their selection of machine operations. Zurich is building a machine of standard components, with drum storage, but attempting primarily to devise an order code that will simplify programming to the utmost. The Ace at NPL, and the forthcoming Deuces, use mercury for the main store and magnetic drum auxiliary, and embody a logic that may have engineering advantages, but appears to complicate the task of programming quite considerably. And the Royal Aircraft Establishment is building a machine for its own use, for purposes of data reduction, with emphasis upon terminal facilities rather more than upon the central computer.

In the use of the machines, their limitations--generally slow speeds and frequently small stores--together with their common use by educational institutions, have encouraged attention to techniques and led to very important developments in numerical analysis. Only at the Bureau of Standards--Washington, and, formerly but no longer, Los Angeles--has there been a group in this country as active as any of several abroad in numerical analysis: Amsterdam, Zurich, Darmstadt, Basel, NPL, Cambridge. On the other hand, except at Cambridge, there seems to be relatively slight interest in what might be called programming theory and in such active development of service routines as can be found at MIT, Remington-Rand, the University of Illinois, and the University of Michigan.

Nevertheless, a partial exception, other than Cambridge, to the latter statement would be the Elliott Brothers. For some time they have had a staff for service computing. Formerly this included ten or a dozen people using desk machines. They have now, however, elaborated their subroutine library to such an extent, that virtually all computations are done electronically and very little use is made of desk machines.

While I feel fairly confident in making the above generalizations comparing European and American activity in numerical analysis, I feel less so in contrasting activity in machine development. In numerical analysis, publications and memoranda provide relatively concrete evidence of activity, and I collected a substantial number of these from the various places visited. The literature on engineering development is much less complete or even representative. A quick, bird's-eye view of foreign activity has provided a certain impression, but I do not have a similar temporal cross-section of domestic activity for a proper comparison. With this caution I shall leave the generalities and consider some of the details.

The International Congress of Mathematicians opened on September 2 and continued through September 9. In conjunction with, but not part of, the Congress proper were three Symposia, of which one, on stochastic processes, opened on September 1, and one, on formal systems, continued over September 10. These sessions I attended. Since I arrived in Amsterdam on August 30, this left August 31 free for visiting the Mathematical Center.

Abstracts of contributed papers at the Congress have already been published, and invited addresses for the Congress and the Symposia are to be published. Hence I shall mention only some conversations, without attempting to summarize any of the papers. One remark on my own paper might be in order. I felt fairly sure that the content, on continued fractions, would be new to most or all the Americans, but I was not so sure about the Europeans. But judging from comments and questions it appeared that continued fractions have not been exploited to any greater extent there than here.

The Mathematical Center at Amsterdam is government supported, and controlled by a board of 6 mathematicians, including van Wijngaarden. Service computing is done for private industry, but mostly the machine is used for research. Personnel, under the direction of the senior mathematicians of the Center, include assistants and computers. The computers are full-time employes whose schooling seems to be roughly equivalent to one year in an American college, with mathematics to, but not including, calculus. The assistants are generally part-time with roughly the equivalent of a master's degree.

Besides an IBM installation, the Center has the ARRA, originally a relay machine, which, however, now has a drum and is being made into an all electronic machine. The Center is making a similar machine for an airplane manufacturer. The drum stores 1024 words of 30 binary digits each. The commands are single-address, with 5 bits for the operation and 10 for the address. Another drum is being added, but since random access to the other drum would require an 11th bit which is not available, the storage on this drum will be fixed, at least for a given problem. As already mentioned, another machine is being built with magnetic core memory, floating decimal representation, and roughly Oracle speeds with completion date approximately two years hence.

On mathematical matters, I learned that van Wijngaarden has been for some time interested in the statistical distribution of generated errors. His work has appeared only in brief summary in an internal memorandum, of which I was able to get a copy for further study. Copies of subsequent publications are promised.

McColl, of Fort Halstead, came by the Center while I was there, and it was then that I learned of the Ferranti machine at ARE. They are using it chiefly for hydrodynamic and aerodynamic calculations.

The sixth Ferranti Computer had just recently been installed at the Shell Laboratories in Amsterdam, except for the drum which had yet to arrive. Interested members of the Congress were taken to see the machine on Monday, September 6. The drum is to store 640,000 bits. Input is, of course, by means of a Ferranti reader, and there are three types of output: a typewriter controlled by the machine directly and giving 6.7 characters (of 5 bits) per second; a punch giving 15 characters per second; and a printer of French make giving 180 lines per minute with a possible 64 characters per line.

My information about the Besk at Stockholm came from G. G. Dahlquist, whom I met during the Congress. The organization operating the machine appears to be roughly comparable to the National Applied Mathematics Laboratories in this country, though on a much smaller scale. It is supported partly by the government, and partly by income from service computing done for industry and other government establishments. The staff is made up entirely of senior mathematicians who carry each problem through completely, including even the coding and machine operation.

Aitken of Edinburgh and Ostrowski of Basel had seen my book on display at the Congress and both offered to supplement my information on some points. Aitken promised to send reprints, and a letter (received on my return to Oak Ridge) describing an iterative method for factoring polynomials with second-order convergence. He did not spell out the proof of convergence, which is laborious, though the method is indicated in the reprints. Ostrowski called my attention to the fact that a theorem I had attributed to Bodewig was in fact due to him, and told me of some of his recent work, too elaborate to be detailed at that time, on iterative solutions of algebraic equations.

My first objective in Zurich was to spend some time with Rutishauser learning about recent developments with the Q D (for quotient-difference) algorithm. His recent paper describing this I consider to be one of the most important contributions made in the field of numerical analysis during recent years. His work subsequent to that already published is mostly contained in two papers to be published about a year hence, and which I was able to borrow and study while I was in Zurich. Without attempting to summarize them here, it may be said that they deal with the following problems: The Bernoulli method of finding solutions of an algebraic equation, and the standard iterative method of finding characteristic roots and vectors of a matrix introduce, as they are generally applied, increasingly large generated error for all roots but the extreme ones. Rutishauser shows how the procedures can be modified so as to evade the difficulty. From the theoretical point of view, these papers also throw interesting additional light on the Q D algorithm itself, and hence upon the other fields with which it is intimately related: continued fraction expansions, summation of divergent series, and the inversion of matrices.

Parenthetically, the practical (but not the theoretical) significance of these papers of Rutishauser's was diminished somewhat by two pieces of information I picked up later. One was that Aitken had published a method as far back as 1931 which achieves essentially the same result for algebraic equations, though the paper is apparently not well known, was probably not known to Rutishauser, and was certainly not known to me. The second was that at NPL a technique due to Wilkinson (who is to visit Oak Ridge in December) enables them to minimize the loss in significance with characteristic roots to a level they have always found tolerable.

Some of the computational problems being worked on at the Technische Hochschule were of interest. One of the most obvious applications of the Q D algorithm is the determination of the periods of harmonic series. A somewhat less obvious application is as follows: The solutions of many differential equations are exactly or approximately representable as sums of exponentials and Murray has given a technique for numerical integration of differential equations which is based upon this fact, but which requires the knowledge of the exponents. Such information is not usually available a priori, but estimates can be made by use of the Q D algorithm.

Leaving the Q D algorithm, Rutishauser showed me the solution of a system of differential equations which had a high degree of instability but which were successfully integrated by means of a technique described by Quade at the Congress. This is simply to differentiate each equation and solve the system of twice the initial order, so that one obtains at each point both the value and the derivative of each dependent variable. According to Rutishauser, the device permits reduction of the number of mesh points by an order of magnitude.

Finally he exhibited a partial differential equation for which a simple iterative scheme failed to converge, whereas relaxation in 6 x 6 blocks converged quite rapidly. It is clear that block relaxation should converge more rapidly than simple relaxation. The significant feature is the conversion of an actually diverging to a converging iteration.

I was somewhat surprised to find that work on programming theory had not progressed beyond earlier publications by Stiefel and Rutishauser except in the formulation of plans for their new machine. The present machine is quite primitive, being slow, relay, and externally programmed. Hence there is perhaps no point in describing it further.

The new machine will have a drum memory of 10,000 words, and in line with the endeavor to simplify programming to the utmost, will be floating decimal. The exponent will vary from -199 to +199, and the mantissa will have 11 decimal digits with the decimal point following the first digit. Commands will be single-address, with two commands per word. Each command contains an I-digit which, when $\neq 0$, signifies that the address of the command is to be modified by the contents of position I. In other words, there are 9 B-tubes. Also each command includes a marking digit called Q which signifies a completed cycle and hence obviates the need for a counter.

Of the three registers, a fixed one, called MR, receives the operand whatever the operation may be, and the same register always exhibits the result. The machine will stop on exponent overflow when operations are floating-point, and on magnitude overflow when operations are fixed-point.

Besides typewriter and tabulator output, there will be punched-card input and output, Remington-Rand having been selected in preference to IBM to permit purchase.

One aspect of computing machine development in England is the use of cold cathode tubes. These have the desirable properties of extremely low heat generation and phenomenally long life (1000 years has been claimed). The disadvantage is that they are slow, operating at kilocycle rather than megacycle rates. At Harwell, in particular, the present machine, and one now under construction, uses them extensively. The present machine is externally controlled, with seven readers for control loops. Each of the

80 storage registers is in effect an accumulator. The speed is such that a skillful operator with a desk machine can keep up with it for short periods. However, it is quite reliable and ordinarily runs overnight unattended.

The new machine, to be finished in December or January, uses point-contact transistors for switching elements and magnetic drum for storage. A test circuit has accumulated several hundred thousand transistor-hours of continuous operation, interrupted with only a single failure, and appears to be quite insensitive to changes in temperature and humidity.

The machine will be binary fixed-point, with 31 digits including sign, but otherwise the logic is extremely unorthodox. Each channel of the drum is divided into 8 arcs, each arc containing 8 interlaced words. One channel is used in effect as a multiple register, and when a word is read into this channel it is repeated on each of the 8 arcs. Thus the register channel stores 8 words, each repeated 8 times, and any number or all of these can be added simultaneously. In any stage of the addition the adder takes the corresponding digits of all addends, adds, shifts, and stores the spilled-off digit. The rest stays as a "carries" to combine with the digits next to the left in the addends. The entire operation requires about 2 ms. Multiplication requires only about three addition times. Unfortunately there is a lack of symmetry in shifting, so that division is made more difficult and may not be included as a special operation.

Professor Williams was not available at the University of Manchester but I spent an afternoon with Kilbourne, discussing machines, and with Brooker, discussing computation. Manchester is the birthplace of the Williams tube for electrostatic storage, and I was interested first in learning about any recent developmental work on these. It turned out, however, that existing tubes meet all their requirements and no such work has gone on for the past two years or so. Instead they are now interested in static reading of magnetic tape, and, as mentioned above, in a transistor computer.

The old Ferranti machine operates 24 hours a day except for 3 to 4 hours off for preventive maintenance. Some work is done for industry, at a rate of £20 per machine hour, and the customers do their own programming. A small programming group under Brooker gives instruction and assistance in programming, and engages in some research on programming and numerical analysis.

Besides the transistor machine being built for engineering research, another machine is being built for computing. This will be floating binary, with parallel electrostatic storage but serial arithmetic. The advantage of this scheme is not at all clear since they have to introduce delay circuits to stagger the bigits as they come out of storage to the registers, and to synchronize them when they go in the reverse direction.

The bigits are arranged in decades, four to a word and two to a command, one for the exponent when the number is floating point. Although additions take 120 μs, multiplications require only 360. A drum provides auxiliary storage. The machine has actually operated, but is not yet regarded as complete.

In many respects the most interesting establishment was that at Lyons Tea Company. There they became interested in possible commercial applications of electronic computers as far back as 1947. Mr. T. R. Thompson came to the U. S. at that time, though there was then very little that he could actually see. Nevertheless, they decided the outlook was promising, so they had some of their engineers assist in the construction of the Edsac at Cambridge, after which they returned to Lyons and built Leo largely following the same design. At the same time a programming group was investigating their real needs and attempting to formulate these in machine terms.

At the present time they use the machine for the analysis of results of market surveys, and for computing the payroll for a sample of 2,500 of their employes. The sample payroll is being run as an experiment, but they now consider that it demonstrates the utility of the operation, intending to include all employes (about 30,000) when they have finished a second machine now being constructed. Other uses, such as stores and inventories, are being studied but they feel that this can be made feasible only after they have devised a technique for direct reading of an initial recording.

The most impressive aspect of the payroll operation was the fact that it was done without auxiliary storage. However, Leo does have a buffer storage unit, a part of the fast memory, into which it is possible to read while computing continues. Effective use is made of multiple inputs, two card readers and two tape readers. Fixed information and variable information are stored separately and read in alternately, an employe at a time. Payments (including the number of coins and notes of each denomination required), deductions, etc., are computed and printed out by means of a tabulator. Input, computation, including numerous cross checks, and printout, require altogether about 1-1/2 sec per employe.

Several references have been made already to Elliott Brothers and the 402, and there is not much more to be said. They became interested in computers as a result of having been commissioned by the Admiralty to look into digital solution of the fire control problem during the war. They later built a special purpose digital computer for the Admiralty, using nickel-acoustic-line and magnetic-disk storage, with the program stored apart from the data.

For service computing they use Nicholas, with 1024 words of 32 digits each, and no auxiliary memory. They have also what is essentially a complete 402 except for the packaging, and this has been run for short periods for test purposes.

At the Royal Aircraft Establishment the greatest volume of work is in data reduction, and hence they have given a great deal of attention to terminal facilities, including analogue-to-digital converters. Among other developments of their own, they have a light airborne paper-tape punch with a speed of 75 characters per second, and a fast punched-card reader, photo-electric, which pushes cards through at a rate of 10 per second and scans at a rate of 1800 columns per second.

Their present computing facilities include a Hollerith installation, a Philbrick differential analyzer, a polynomial equation solver, and a linear equation solver. The last item is intended to solve systems of order 12 or less, analogue, to three decimal digits, and for the system

$Ax = y$ it is designed to do a Seidel iteration on $A^T Ax = A^T y$. This being the case, it gives a least-squares solution of an over-determined system of not more than 12 equations. Unfortunately, the operating record has been very poor, and it was not operating when I was there.

The polynomial equation solver, also analogue, was much more successful. An equation of degree 6 or less is written in the form $f(z) \equiv a_3 z^3 + \dots + a_{-3} z^{-3} = 0$, and although the inputs make provision for only real a's, the internal circuitry could handle complex a's as well. The oscilloscope then exhibits the curve $f(r\,e^{i\theta})$ for fixed r. It is a simple matter to vary r and find those values of r and of θ at which the curve passes through the origin.

Pending completion of their copy of the Deuce, they use time on the Ace, with their own staff to program and operate. The Rascal is being built to their own design, with floating-point decimal, and a cyclic progressive binary-coded representation of the decimal integers which is intended to minimize the effects of punching and reading errors. (It is not clear that the effect will be appreciable. Whereas consecutive digits differ by only one hole, it is a particular hole, and an error in one hole will in general lead to an error greater than 1.)

The Mathematical Division of the National Physical Laboratory corresponds roughly to the National Applied Mathematics Laboratories of the National Bureau of Standards. They sell time on the Ace and will sell time on the Deuce, and they do service computing and programming, being thus partially self-supporting. However, they are in no way dependent upon this income and hence have considerably more freedom for basic research. This is, in fact, one of the major centers for research on numerical analysis. Their operating philosophy, and their computational experience, were therefore of considerable interest.

While their library is reasonably complete on functional evaluation, they do not make use of special input and compiler routines. Instead, each problem is coded directly in machine form. Moreover, Goodwin was emphatic on their principle of avoiding floating point if at all possible. This general philosophy seems to be in marked contrast to that of the Elliott group and doubtless accounts in part for the fact that NPL maintains a considerable staff for hand computing whereas Elliott's does not.

Many of their service computations have been on classified problems, of which they knew only the mathematical formulation. Moreover, these have not included partial differential equations since, until just recently, they have had no auxiliary storage (they now have a drum). They have had a considerable number of matrix inversions, of latent root and vector computations, and, sporadically, algebraic equations.

On matrix inversion, they use straight elimination and handle symmetric and nonsymmetric matrices alike. They search for the largest element in a column, and program a stop in case of overflow. In principle there could be one overflow in any column. Nevertheless the stop has never occurred. In general they do not believe in trying to provide for contingences that may very well not arise.

Their treatment of latent roots and vectors, already mentioned, was the most remarkable point brought up. They use straight iteration to get the largest root and the associated vector; then they reduce the order of the matrix by one, and proceed. In case of slow convergence they apply

Aitken's δ^2-process once, and then continue the iteration. They have handled "hundreds" of matrices, both symmetric and otherwise, of orders up to 30, and claim greater accuracy even down to the last root than they can get by any other method they have tried, including that of Lanczos.

In principle this method can be adapted for algebraic equations, although apparently they have not done so. Instead Olver recommends either using Graeffe's method for exploration, or else digitalizing essentially the process used by the RAE analogue equation solver, and then following up by Newton's method for each root.

In a paper by Goodwin and Fox on integral equations, they had passed over equations of the first kind rather hurriedly, and I was hoping they might have some techniques they had withheld. Unfortunately they did not.

The Edsac II at Cambridge has already been mentioned and briefly characterized: magnetic cores for storage and switching; permanently stored input routines and standard subroutines in supplementary storage; roughly Oracle speeds; systematic use of micro-programming for the construction of basic operations; two B-tubes; both fixed- and floating-point operations; availability of words and half-words.

On Edsac I they have recently installed a drum, and added a B-tube. For decimal punch-out they use a two-out-of-five code, which has checking advantages but requires a special punch-out routine. Otherwise the machine is essentially as described in the published literature.