

to the dash period which is 4 microseconds followed by a blackout period, during movement to the next digit, a total of 10 microseconds. Speed of operation has not been a prime object of design. We have been content to accept the natural speed of operation of the components used in the machine and we have taken full advantage of the fact that we knew we were going to have a large storage capacity in the machine, giving quite a good speed of operation. Automatic control is applied only to the gain of the pick-up amplifier.

G. E. Reynolds (United States Air Force, Cambridge Research Center): I have a

question for Dr. Kilburn. The random number generator that you mentioned—may we have a few details?

T. Kilburn: It operates on the principal that one takes a noise source and allows it to control a counter. The state of the counter is inspected every so often and the nought or one obtained is put in the accumulator of the machine and moved upwards before the next nought or one is put in. This continues until a 20-digit number is assembled in the accumulator.

E. Blumenthal (Eckert-Mauchly Computer Corporation): I would like to ask of either speaker the reason for using a servo

system for the master oscillator to control the drums rather than the sprockets and using that as your oscillator.

T. Kilburn: This is an opening sentence in a debate that never ends, in England. I will tell you why we did that, and you can be the judge as to whether it is right or wrong.

In the early stages, we did not know whether or not we would have more than one drum inside the machine. If we did require this, then we should have had to synchronize the second drum to the first. Thus we faced this problem of locking in the drum to a master oscillator immediately. The servo-mechanism is extremely reliable.

The Whirlwind I Computer

R. R. EVERETT

PROJECT Whirlwind is a high-speed computer activity sponsored at the Digital Computer Laboratory, formerly a part of the Servomechanisms Laboratory, of the Massachusetts Institute of Technology (M.I.T.) by the Office of Naval Research (O.N.R.) and the United States Air Force. The project began in 1945 with the assignment of building a high-quality real-time aircraft simulator. Historically, the project has always been primarily interested in the fields of real-time simulation and control; but since about the beginning of 1947 most of its efforts have been devoted to the design and construction of the digital computer known as Whirlwind I (WWI). This computer has been in operation for about 1 year and an increasing proportion of project effort now is going into application studies.

Applications for digital computers are found in many branches of science, engineering, and business. Although any modern general-purpose digital computer can be applied to all these fields, a machine is generally designed to be most suited to some particular area. Whirlwind I was designed for use in control and simulation work such as air traffic control, industrial process control, and aircraft simulation. This does not mean that Whirlwind will not be used on applications other than control. About one-half the available computing time for the next year will be

assigned to engineering and scientific calculation including research in such uses supported by the O.N.R. through the M.I.T. Committee on Machine Methods for Computation.

These control and simulation problems result in a specialized emphasis on computer design.

SHORT REGISTER LENGTH

WWI has 16 binary digits and the control problems are usually very simple mathematically. Furthermore, the computer is almost always part of a feedback rather than an open-ended system. Consequently, roundoff errors are seldom troublesome and the register length can be shortened to something comparable to the sensitivity of the physical quantities involved, perhaps five decimal places or less.

WWI has a register length of 16 binary digits including sign or about four and one-half decimals. The register length was chosen as the minimum that would provide a usable single-address order, in this case five binary digits for instruction and 11 binary digits for address. In a future machine we would probably increase this register length to 20 or 24 binary digits to get additional order flexibility; the increased numerical precision is less important.

For scientific and engineering calculation, greater than 16-digit precision is often required. There is available a set of multiple-length and floating point sub-

routines which make the use of greater precision very easy. It is true that these subroutines are slow, bringing effective machine speed down to about that obtained by acoustic memory machines. It is much more efficient occasionally to waste computing time this way than continuously to waste a large part of the storage and computing equipment of the machine by providing an unnecessarily long register.

HIGH OPERATING SPEED

WWI performs 20,000 single-address operations per second. Control and simulation problems require very high speeds. The necessary calculations must be carried out in real time; the more complex the controlled system is, the faster the computer must be. There is no practical upper limit to the computing speed that could be used if available.

Where the problems are large enough, and these problems are, one high-speed machine is much better than two simpler machines of half the speed. Communication between machines presents many of the same problems that communication between human beings presents.

Great effort was put into WWI to obtain high speed. The target speed was 50,000 single-address operations per second, and all parts of the machine except storage meet this requirement. The actual WWI present operating speed of 20,000 single-address operations per second is on the lower edge of the desired speed range.

LARGE INTERNAL STORAGE

WWI now has 1,280 registers. A large amount of high-speed internal storage is needed since it is not in general possible to use slow auxiliary storage because of

R. R. EVERETT is at the Digital Computer Laboratory, Massachusetts Institute of Technology, Cambridge, Mass.




Figure 1. Sample computer output

the time factor. In many cases a magnetic drum can be useful since its access time is short compared to the response times of real systems. Even with a drum there is considerable loss of computing and programming efficiency due to shuffling information back and forth between drum and computer.

WWI is designed for 2,048 registers of storage. Until recently there has been available only about 300 registers. This number, while small, has been adequate for much useful work. Very recently a second bank of new-model storage tubes has been added. These new tubes operate at 1,024 spots per tube bringing the total WWI storage to 1,280 registers. These tubes have been in the computer and under test for 2 months and in active use for about 2 weeks. In the next few months the tubes in the first bank will be replaced by new model storage tubes bringing the total storage to 2,048. This number is on the lower end of what the project considers desirable. What the computer business needs, has needed, and will probably always need is a bigger, better, and faster storage device.

EXTREME RELIABILITY

In a system where much valuable property and perhaps many human lives are dependent on the proper operation of the

computing equipment, failures must be very rare. Furthermore, checking alone, however complete, is inadequate. It is not enough merely to know that the equipment has made an error. It is very unlikely that a man, presumably not too well suited to the work during normal conditions, can handle the situation in an emergency. Multiple machines with majority rule seem to be the best answer. Self-correcting machines are a possibility but appear to be too complicated to compete, especially as they provide no stand-by protection.

The characteristics of the Whirlwind I computer may be recapitulated as follows:

Register length	16 binary digits, parallel
Speed	20,000 single-address operations per second
Storage capacity	Originally 256 registers Recently 320 registers Presently 1,280 registers Target 2048 registers
Order type	Single-address, one order per word
Numbers	Fixed point, 9's complement
Basic pulse repetition frequency	1 megacycle
Tube count	2 megacycles, (Arithmetic element only)
Crystal count	5,000, mostly single pentodes
	11,000

There are 32 possible operations, of which about 27 are assigned. They are of the usual types; addition, subtraction, multiplication, division, shifting by an arbitrary number of columns, transfer of all or parts of words, subprogram, and conditional subprogram. There are terminal equipment control orders and there are some special orders for facilitating double-length and floating-point operations.

One way to increase the effective speed of a machine is to provide built-in facilities for operations that occur frequently in the problems of interest. An example is an automatic co-ordinate transformation order. The addition of such facilities does not affect the general-purpose nature

of the machine. The machine retains its old flexibility but becomes faster and more suited to a certain class of problems.

From March 14, 1951, at which time we began to keep detailed records, until November 22, 1951 a total of 950 hours of computer time were scheduled for applications use. The machine has been running on two shifts or a total of about 3,000 hours during this interval. The two-thirds time not used for applications has been used for machine improvement, adding equipment, and preventive maintenance.

Of the 950 hours available, 500 have been used by the scientific and engineering calculation group, the rest for control studies. The limited storage available until recently has been admittedly a serious handicap to the scientific and engineering applications people. There has not been room in storage for the lengthy subroutines necessary for convenient use of the machine. The largest part of their time has been spent in training, in setting up procedures, and in preparing a library of subroutines.

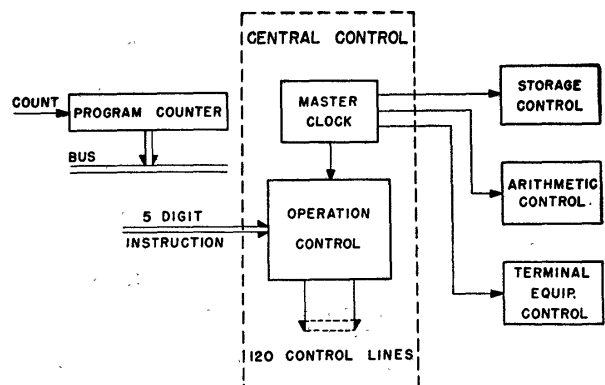
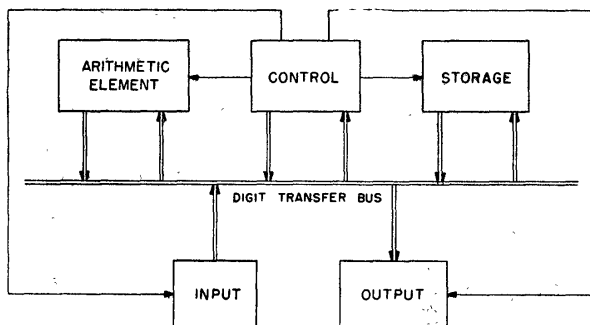
A partial list of the actual problems carried out by the group includes:

1. An industrial production problem for the Harvard Economics School.
2. Magnetic flux density study for our magnetic storage work.
3. Oil reservoir depletion studies.
4. Ultra-high frequency television channel allocation investigation for Dumont.
5. Optical constants of thin metal films.
6. Computation of autocorrelation coefficients.
7. Tape generation for a digitally-controlled milling machine.

The scientific and engineering applications time on Whirlwind I has been organized in a manner patterned after that originated by Dr. Wilkes at EDSAC. The group of programmers and mathematicians assigned to WWI assist users in setting up their own problems. Small problems requiring only a few seconds or

Figure 2 (below). Simplified computer block diagram

Figure 3 (right). Control



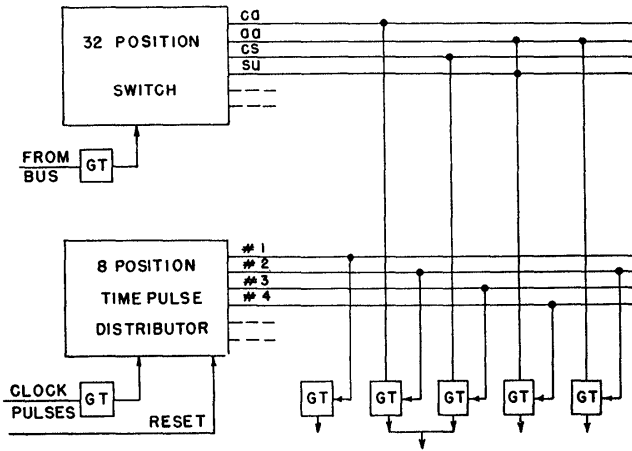


Figure 4. Operation control

type. Unfortunately, the simple concept often becomes complicated in execution, and this is true here. WW's control has been complicated by the decision to keep it completely flexible; the arithmetic element by the need for high speed, the storage by the use of electrostatic storage tubes, the terminal equipment by the diversity of input and output media needed.

Control

The WW control is divided into several parts, as shown in Figure 3.

CENTRAL CONTROL

The central control of the machine is the master source of control-pulses. When necessary the central control allows one of the other controls to function. In general there is no overlapping of control operation; except for terminal equipment control, only one of the controls is in operation at any one time.

STORAGE CONTROL

Storage control generates the sequence of pulses and gates that operate the storage tubes. Central control instructs the storage control either to read or to write.

ARITHMETIC CONTROL

Arithmetic control carries out the details of the more complex arithmetic operations such as multiplication and division. The setup of these operations plus the

minutes of computer time are encouraged. Applications time is assigned in 1-hour pieces two or three times a day. No program debugging is allowed on the machine. Program errors are deduced by the programmer from printed lists of results, storage contents, or order sequences as previously requested from the machine operator. The programmer then corrects his program which is rerun for him within a day or perhaps within a few hours.

Every effort is made to reduce the time-consuming job of printing tabulated results. In many cases a user desires large amounts of tabulated data only because he doesn't really know what answers he wants and so asks for everything. Such users are encouraged to ask only for pertinent results in the form of numbers or

curves plotted by the machine on a cathode-ray tube and automatically photographed. If these results prove inadequate or the user gets a better idea of his needs, he is allowed to rerun his program, again asking only for what appear to be significant results. Figure 1 shows a sample curve plotted by the computing machine showing calibrated axes and decimal intercepts.

WWI System Layout

Figure 2 shows the major parts of any computer such as WWI. The major elements of the computer communicate with each other via a central bus system.

WWI is basically a simple, straightforward, standard machine of the all-parallel




Figure 5 (left). View of central control




Figure 6 (right). View of electrostatic storage

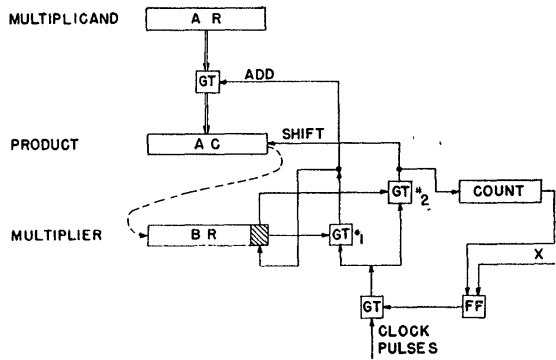



Figure 7 (above). Arithmetic element

Figure 9 (right). View of test control



complete controlling of the simpler operations such as addition are carried out by central control.

TERMINAL EQUIPMENT CONTROL

Terminal equipment control generates the necessary control pulses, delay times, and interlocks for the various terminal equipment units.

PROGRAM COUNTER

The program counter which keeps track of the address of the next order to be carried out is considered as part of control. This is an 11-stage binary counter with provision for reading to the bus.

Most of the functions of these subsidiary controls could be combined with the central control. The major reason they are not is that they were designed at different times. The arithmetic element and its control came first, followed by central control. At the time central control was designed, the necessary characteristics of storage control were unknown. In fact, the machine was designed so that any parallel high-speed storage could be used. The form of terminal equipment control was also unknown at this time. Since flexibility was a prime specification, it was felt preferable to build separate flexible controls for the various parts of the computer than to try to combine all the needed flexibility in one central control.

In a new machine we would attempt to combine control functions where possible, hoping to have enough prior knowledge about component needs to eliminate subsidiary controls completely. We would still insist on a large degree of control flexibility.

MASTER CLOCK

The master clock consists of an oscillator, pulse shaper and divider that generate 1- and 2-megacycle clock pulses, and a clock pulse control that distributes

these clock pulses to the various controls in the machine. It is this unit that determines which of the subsidiary controls actually is controlling the machine. This unit also stops and starts the machine and provides for push-button operation.

Operation Control

The operation control, see Figure 4, was designed for maximum flexibility and minimum number of operation digits, and, consequently, minimum register length. It is of the completely decoding type.

The operation switch is a 32-position crystal matrix switch that receives the 5-bit instruction from the bus and in turn selects one of 32 output lines corresponding to the 32 built-in operations.

There are 120 gate tubes on the output of the operation control. Pulses on the 120 output lines go to the gate drivers, pulse drivers, and control flip-flops all over the machine; 120 is a generous num-

ber. The suppressors of these gate tubes are connected to vertical wires that cross the 32 output lines from the operation switch. Crystals are inserted at the desired junctions to turn on those gate tubes that are to be used for any operation.

The time pulse distributor consists of an 8-position switch driven from a three binary-digit counter. Clock pulses at the input are distributed in sequence on the eight output lines. The control grids of the output gate tubes are connected to these timing lines. The output of the operation control is thus 120 control lines on each of which can appear a sequence of pulses for any combination of orders at any combination of times.

Central Control


The Central Control of the machine is shown in Figure 5. The control switch is in the foreground with the operation matrix to the right.

Electrostatic Storage

The electrostatic storage shown in Figure 6 consists of two banks of 16 storage tubes each. There is a pair of 32-position decoders set up by address digits read in from the bus. There is a storage control that generates the sequence of pulses needed to operate the gate generators, et cetera. A radio frequency pulser generates a high power 10-megacycle pulse for readout.

Each digit column contains, besides the storage tubes, write plus and write minus gate generators and a signal plate gate generator for each tube. Ten-megacycle grid pulses are used for readout in order to get the required discrimination between the fractional volt readout pulses and the 100-volt signal plate gates. For each storage tube there is a 10-megacycle amplifier, phase-sensitive detector and gate

Figure 8. View of arithmetic element



tube, feeding into the program register. The program register is used for communicating with the storage tubes. Information read out of the tubes appears in the program register. Information to be written into the tubes must be placed in the program register.

Arithmetic Element

The arithmetic element, see Figure 7, consists of three registers, a counter, and a control.

The first register is an accumulator (*AC*) which actually consists of a partial-sum or adding register and a carry register. The accumulator holds the product during multiplication.

The second or *A*-register holds the multiplicand during multiplication. All numbers entering the arithmetic element do so through *AR*.

The third or *B*-register holds the multiplier during multiplication. The accumulator and *B*-register shift right or left. A high-speed carry is provided for addition. Subtraction is by 9's complement and end-around-carry. Multiplication is by successive additions, division by successive subtractions, and shift orders provide for shifting right or left by an arbitrary number of steps, with or without roundoff.

The arithmetic element is straightforward except for a few special orders and the high speed at which it operates. Addition takes 3 microseconds complete with carry; multiplication, 16 microseconds average including sign correction.

In figure 8 are shown several digits of the arithmetic element. The large panels are accumulator digits. Above the accumulator is the *B*-register below it the *A*-register.

Test Control

Test control, shown in Figure 9, is used at present both for operating and for trouble shooting the computer. The control includes:

1. Power supply control and meters.
2. Neon indicators for all flip-flops in the machine.
3. Switches for setting up special conditions.
4. Manual intervention switches.
5. Oscilloscopes for viewing wave forms. A probe and amplifier system allows viewing any wave form in the computer on one scope at test control.
6. Test equipment to provide synchronizing, stop, or delay pulses at any step of any

order of a program, allowing viewing wave forms on the fly anywhere in the machine.

An important part of the test facilities is the test storage, a group of 32 toggle-switch registers plus five flip-flop registers that can be inserted in place of any five of the toggle-switch registers. This storage has proved invaluable not only for testing control and arithmetic element before electrostatic storage was available but also for testing electrostatic storage itself. When not in use for test purposes test storage earns its keep as part of the terminal equipment system. The toggle-switches hold a standard read-in program; the flip-flop registers are used as in-out registers for special purposes.

CHECKING

Logical checking facilities built into WWI are rather inconsistent. A complete bus transfer checking system has been provided, duplicate checking of some terminal equipment is permitted, but little else is thoroughly checked. We felt that it was worth while to thoroughly check some substantial portion of the machine. This portion would then serve as a prototype for studying the tube circuitry used throughout the machine. We did not feel it was worth while to check all the machine, a procedure that requires a great deal of added equipment and logical complexity plus a substantial loss in computing speed.

Operating experience has shown us that it is not worth while to provide detailed logical checking of a machine. In a new machine we would leave out the transfer checking. The amount of information and security given by the detailed checking system is not enough to warrant the expense of building and maintaining it.

This decision is based on the expectation that a computing machine should operate 95 per cent of total time or better and that the average time between random failures should be of the order of 5 to 10 hours or approximately 10^9 operations.

In our opinion the way to achieve the extremely high reliability needed in some real-time control problems is to provide three or more identical but distinct machines, thus obtaining error correction as well as detection, plus such features as standby, safety, and damage control. Even so the failure probability of each machine must be kept low by proper design, marginal checking, and preventive maintenance.

Extremely high reliability means a reliability far beyond that achieved in existing machines and not conveniently rep-

resented as a per cent. Consider a system consisting of three machines, each operable 98 per cent of the time and each averaging 10 hours between random errors.

One machine will be out of operation 1/2 hour per day.

Two machines will be out of operation 1/4 hour per month.

All three machines will be out of operation 4 minutes per year. Furthermore undetected random errors might occur on the average of once a year. Such reliability is needed in some systems.

Our decision to omit detailed checking does not extend to checking devices intended to detect programming errors. Devices to check for overflow from the arithmetic element or for nonexistent order configurations are necessary. Programmers make many mistakes. Techniques for dealing with programming errors are very important and need future development.

TERMINAL EQUIPMENT

At the present time, Whirlwind is using the following terminal equipment:

1. A photoelectric paper tape reader.
2. Mechanical paper tape readers and punches.
3. Mechanical typewriters.
4. Oscilloscope displays 5 to 16 inches in diameter with phosphors of various persistencies including a computer-controlled scope camera.
5. Inputs from various analogue equipments needed for control studies.
6. Outputs to analogue equipment.

To be added during the next year:

1. Magnetic Tape (units by Raytheon). One such unit is now being integrated with machine.
2. Magnetic drums (units by Engineering Research Associates, Inc.).
3. Many more analogue inputs and outputs.

This great complexity of terminal equipment requires a flexible switching system. There is a single in-out register (*IOR*) through which most of the data passes.

There is a switch which is set up by an order to select the desired piece of terminal equipment. Other orders put data into *IOR* or remove data from *IOR*. The in-out control provides the necessary control pulses to go with each type of equipment. In general the computer continues to run during terminal equipment wait times; suitable interlocks are provided to prevent trouble. This complete equipment has not yet been fully installed.

(Discussion of this paper was combined with that of the following paper)

Evaluation of the Engineering Aspects of Whirlwind I

NORMAN H. TAYLOR

THE Whirlwind system is an expanding one now adding terminal equipment such as magnetic tapes, drums, and scopes. The core of the system including control, memory, and arithmetic element now has been running long enough to make an engineering evaluation of its performance, so the following comments will apply mainly to these three sections.

Measured in terms of applications performance, Whirlwind has been used about 950 hours in the last 6 months, but the basic elements have accumulated over 8,000 hours of use during engineering tests over the last year and a half.

A current running day in the Whirlwind system has 5 to 8 hours assigned to application engineering, 6 hours to engineering development such as the extension of terminal facilities and the increase in the internal storage capacity, and 1 hour to marginal checking.

To express the performance record during the applications period, we have taken a ratio of the total number of productive hours to the total assigned hours and get a figure of 85 per cent as our over-all efficiency. A more detailed study of the remaining 15 per cent of time indicates that the people using the machine are responsible for about two-thirds of this down time while the machine itself is responsible for about one-third. This is due to the fact that the machine schedules are made 1 week in advance and if an engineering test takes an hour longer than the schedule indicates, this is counted as down time. Omitting these scheduling problems, machine efficiency is close to 95 per cent.

Another way of evaluating the system is to look at the number of replacements of components in a period of time. Whirlwind now contains 4,450 tubes, most of which have been in their sockets about 8,000 hours. Of these, 616 have been removed for a total of 14 per cent tube removal. This gives a failure rate of 1.7 per cent per 1,000 hours.

Whirlwind contains about 11,000 crystals in two rather widely separate applications. A large group of crystals, over

9,000, are used in low forward current and low back voltage applications and the performance record has been quite satisfactory. Only 180 replacements have been necessary in 8,000 hours for a total of 2 per cent. Expressed in the failure rate per 1,000 hours this number is 0.25 per cent.

A smaller group of 1,048 crystals has been used in clamping circuits at rated back voltage but much poorer results were attained, and 334 units or 31 per cent of these crystals have been removed for a rate of 4 per cent per 1,000 hours.

Before looking into the details of just where these failures arose, it may be well to say that the nature of a failure is very often more important than the fact that a failure exists; for instance: the marginal checking procedure which is an integral part of the Whirlwind system, makes it very easy for one to locate a slowly deteriorating component before it causes a systems failure, and replacement of such a component during a period of marginal checking or preventive maintenance does not hinder the user of the computer. A sudden failure, however, due to an intermittent joint can usually not be predicted. Such failures are apt to occur during an important calculation and breakdown at this time can be very expensive and time consuming. Over half of the failures in the Whirlwind system are now being picked up during the marginal checking period and these do not represent a major obstacle to its use in application periods. The effectiveness of this marginal checking lies somewhere between 50 and 70 per cent of the total failures which do occur. Our experience indicates that one may expect 6 to 10 hours of trouble-free operation after each marginal checking period. On the average, three tubes or crystals are replaced during the daily preventive maintenance period, but a breakdown of the system during an applications period only occurs about once in a 3-day period.

Breakdown of Performance Numbers

The numbers just quoted represent over-all systems performance. To discuss which portions are good and which are in-

ferior from an engineering point of view, the system has been separated into sections as follows:

1. The flip-flop and gate used in arithmetic circuits and portions of control.
2. The control selection system.
3. The memory system (electrostatic storage tube circuitry).
4. Terminal equipment.

To evaluate each of these sections of Whirlwind, a look at the replacement rate of components is enlightening. This rate is to be based on replacements of vacuum tubes and germanium diodes because troubles can usually be traced to failure of one of these two components. Transformers, condensers, delay lines, and resistors are considerably better than tubes or crystals and may be considered quite adequate for use in computers of the future.

First, the gate circuit: A coincident *and* gate is made up of a 7AK7 gate tube, a transformer, and diode as shown in Figure 1. The grids are usually held at -15 volts, and a 20-volt pulse 0.1 microsecond wide is applied to grid number one. When the gate is sensed in this way, if grid 3 is at -15 volts no plate current flows; if grid 3 is at ground, about 30 milliamperes of plate current flows and a 0.1 microsecond pulse passes to the next circuit.

There are 1,300 gate circuits such as this in Whirlwind I (WWI), and the over-all failure rate of these circuits is the very best in the system. It is less than 0.5 per cent per 1,000 hours. The rate of failure seems to be linear, so if we dare to extrapolate, this slope indicates that half of our gate tubes will still be in use at 100,000 hours. Almost all the failures in this circuit are intermittent or mechanical in nature. Improvement in performance would undoubtedly demand considerable improvement in the method of production of vacuum tubes with special attention to cathode flaking, grid spacing, welding, and vacuum techniques.

The excellent performance of this circuit is due, of course, to the availability of the 7AK7 gate tube and its rather ideal application. This tube was designed and built for WW by Sylvania Electric Products, Inc. It is made in a pilot plant with special attention paid to cleanliness, uniformity of parts, vacuum pumping, aging, and choice of cathode materials. Our experience with other tubes indicates that the removal of this tube from the pilot plant production could easily reduce its performance to the level of other tubes.

The flip-flop circuit is a second basic element in the Whirlwind system and is

NORMAN H. TAYLOR is at the Digital Computer Laboratory, Massachusetts Institute of Technology, Cambridge, Mass.