

NOV 29 1982

RJ3117 (38413) 4/10/81
Computer Science

Research Report

HISTORY OF THE DESIGN OF THE SAGE COMPUTER - THE AN/FSQ-7

Morton M. Astrahan
IBM Research Laboratory
San Jose, California 95193

John F. Jacobs
MITRE Corporation
Bedford, Massachusetts 01730

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties).

IBM

Research Division
Yorktown Heights, New York · San Jose, California · Zurich, Switzerland

Copies may be requested from:
IBM Thomas J. Watson Research Center
Distribution Services
Post Office Box 218
Yorktown Heights, New York 10598

RJ3117 (38413) 4/10/81
Computer Science

History of the Design of the SAGE Computer- the AN/FSQ-7

Morton M. Astrahan

John F. Jacobs

IBM Research Laboratory
San Jose, California 95193

MITRE Corp.
Bedford, Massachusetts 01730

ABSTRACT:

This is the story of the development of the SAGE (Semi-Automatic Ground Environment) Air Defense Computer, the AN/FSQ-7. At the time of its operational deployment beginning in 1958, the AN/FSQ-7 was the first large-scale, real-time digital control computer supporting a major military mission. The AN/FSQ-7 design, including its architecture, components and computer programs, drew on R&D programs throughout the United States, but it drew mostly on work being done at MIT Project Whirlwind and at IBM. How all this came about is the subject of this paper.

INTRODUCTION

This is the story of the development of the SAGE (Semi-Automatic Ground Environment) Air Defense Computer, the AN/FSQ-7. At the time of its operational deployment beginning in 1958, the AN/FSQ-7 was the first large-scale, real-time digital control computer supporting a major military mission. Its development was initiated at a time when the perception among Department of Defense (DoD) officials was that Soviet bombers carrying nuclear bombs were a primary threat to the United States. The generally-held belief in the validity of this threat gave the SAGE program the highest DoD priority. The AN/FSQ-7 design, including its architecture, components and computer programs, drew on R&D programs throughout the United States, but it drew mostly on work being done at MIT Project Whirlwind and at IBM. How all this came about is the subject of this paper.

SAGE system programming is an interesting story in its own right, but is outside the scope of this paper. Similarly, the system for management of deployment worked out among the Air Defense Engineering Services project office, Lincoln, the Air Defense Command and the contractors deserves a more thorough treatment.

PROLOGUE

The need for air defense was driven home in the United States in 1941 by the Japanese with their attack on Pearl Harbor. Pearl Harbor demonstrated the need for air surveillance, warning, and real-time control. Shaken by Pearl Harbor, the United States became serious about air defense within its continental limits. By the end of World War II, there were more than 70 Ground Control of Intercept (GCI) sites.

Each GCI site consisted of one or two search radars, a height finder radar, and ground-to-air and air-to-ground communications. The operators sat in front of plan position indicators (PPI), which presented the air situation on a scope with long-persistence phosphors. Aircraft appeared as "blips" of light on the face of the tube. Information on targets from adjacent sites was cross-told by voice telephone. The control centers were usually built around a large, edge-lit plexiglas board which showed the local geographic features. Aircraft of interest were marked on the board by operators standing on scaffolding behind the board and using grease pencils. The big board also showed status information, which was written backwards by the operators. The network of GCI sites became known later as the "Manual System".

Following the Allied victory, the most powerful air forces were in the hands of the Allies, including Russia. There seemed no justification for the expense of maintaining the radar sites established during the war, and support eroded.

In 1947 the Army Air Force was organized as a separate service, reporting

to a newly established Defense Department. The Air Force was given the air defense mission and proceeded to plan the revival of the Manual System. The importance of this mission was increased with the subsequent Russian achievement of producing nuclear bombs, and was further strengthened by later events in Korea. Meanwhile, the Air Force Chief of Staff, Gen. Hoyt S. Vandenberg, became more and more concerned about United States vulnerability to airborne attack. The Air Force Scientific Advisory Board was exposed to the problem, and in 1949, the Board set up an Air Defense Systems Engineering Committee (ADSEC) under George E. Valley, a physics professor at MIT. The Committee became known as the Valley Committee.

The Valley Committee began by looking at the newly reactivated air defense system. This system had been authorized by Congress through the Air Force, and consisted of about 70 Ground Control Intercept radar sites. Except for improved radars and height finders, it was quite similar to the Manual System air defense setup established during World War II. The Committee quickly concluded that the air defense system as reshaped by the Air Force had very low capability, and it recommended that a competent technical organization look into what could be done to improve the system in the short run. As a result, the Western Electric Company and the Bell Telephone Laboratories were given the task of upgrading the existing system; this was to become the Continental Air Defense System (CADS) project. The Valley Committee also suggested that a longer-range look be taken at the problem. It recommended the extensive use of automation, particularly computers, to handle the bookkeeping, surveillance and control problems in the implementation of next generation air defense systems. This conclusion was supported by the development of the Whirlwind computer at MIT. The Whirlwind promised to provide real-time control over a large number of aircraft. It was also noted that the ability to pass digital information over phone lines had been demonstrated at Bell Telephone Laboratories and at the Air Force Cambridge Research Laboratory. To deal with one of the major problems, low altitude surveillance, the Committee recommended the establishment of a large number of short-range low maintenance radars which would be placed closely together to fill gaps in coverage.

The Valley Committee report led General Vandenberg in December of 1950 to ask MIT to establish a laboratory for air defense research and development. The Air Force Scientific Advisory Board endorsed this request and also asked MIT to undertake an interim study of the air defense problem. The study, called Project Charles, ran from February to August of 1951. It gave further support to the concept of a computer-based system. The laboratory was established within MIT in 1951 as Project Lincoln, and in 1952 became the MIT Lincoln Laboratory. The SAGE system evolved from the work of this laboratory [5,9].

PROJECT WHIRLWIND

The Whirlwind computer project at MIT's Digital Computer Laboratory was of crucial importance to the development of the AN/FSQ-7 for several reasons.

First, it provided a demonstration of real-time control by digital computer without which the SAGE project could not have been approved. Second, it provided a reservoir of people with the skills and experience needed to participate in the SAGE system design and development. Third, it provided an experimental testbed for the system design. The story of the Whirlwind project and the role of key people like Jay W. Forrester and Robert R. Everett has been described in [8] and [9].

THE CAPE COD SYSTEM

In the Spring of 1952, the Digital Computer Laboratory (DCL) was working closely with the Lincoln Laboratory and the DCL operations and people concerned with air defense were merged into Lincoln as Division 6. The Whirlwind computer was working well enough to be used as part of Lincoln Lab's experimental air defense system, called the Cape Cod System. This consisted of a control center at the Barta Building in Cambridge, Massachusetts, where Whirlwind was housed, as well as an experimental long-range radar on Cape Cod at South Truro, Massachusetts, and a number of short-range radars called "gap fillers". The control center contained computer controlled operating stations for interaction with human operators. It was equipped with UHF communications to aircraft supplied by the Air Research and Development Command and the Air Defense Command, for the purposes of creating a realistic test of the system.

The Valley Committee and the Charles Study had indicated that a preferred solution for dealing with the low altitude detection problem was to connect together many radars (preferably short-range, low maintenance radars) and make a composite picture of the air situation out of the data taken from these radars. It was largely the need to deal with so much data that had prompted the Valley Committee to favor the use of the computer aids in processing the data in real time. Just as Whirlwind had the potential for filling the needs for this additional data load, work at the Air Force Cambridge Research Laboratory under Jack Harrington on digital transmission of radar data had the potential for filling another need: communicating the data. Harrington's group had developed a technique (actually, several techniques) for transmitting this data. One technique, called slowed-down video, divided the coverage area of short-range radars into a large number of wedge-shaped boxes, the number bounded by range resolution required and the angular resolution that one could achieve with the radar. The boxes were mapped onto a stream of bits sent on a phone line. The stream was synchronized with the radar pulses and the angular position of the radar. Each bit was a '1' if the corresponding box contained a signal return above a certain magnitude; otherwise it was '0'. This technique showed promise for short-range radars, but it was far too inaccurate for the long-range radars.

The Cambridge Research Laboratory (CRL) was also working on methods of providing more angular precision than could be achieved by means of beam forming. One scheme which eventually resulted in another SAGE development, called the AN/FST-2 [7], derived from beam-splitting experiments carried on at CRL. It depended upon the fact that, as a radar beam

retates, the pulse rate is high enough so that several returns are received from a single aircraft. Harrington's group invented a device which determined the center of the target after the beam had swept over it. This device made it possible to increase the angular accuracy by an order of magnitude. Harrington's team also developed a scheme for sending generalized digital data over a standard phone line that had been adapted to the Cape Cod System. Harrington and many of his team from CRL joined Lincoln Laboratory when it was instituted. As soon as the Whirlwind computer was able to perform, an experimental MEW radar at Hanscom Field was connected by phone line to Whirlwind, and the first tracking programs were developed. By 1952, the Cape Cod team had demonstrated the ability of the computer to track and control aircraft in small numbers. The Cape Cod System was intended to demonstrate the operations that were to be executed for field use, in particular the surveillance function and weapons control function. Both of these functions required information on the position of hostile, friendly and neutral aircraft. Some scheme wherein all of the operators in the control center worked from the same positional data base became a requirement.

In the scheme that was adopted, target data was transmitted to the Center in angular coordinates. There the computer translated it into Cartesian coordinates and combined it with the position of the radar that picked up the data, so that each piece of data had an X-Y position in a common coordinate system and could be compared with stored track data (successive positions of an object being tracked). Each operating station was equipped with a console which had a cathode ray tube situation display which combined track and map data. During the course of the operating cycle, the computer presented successive data locations to an X-Y deflection register which simultaneously positioned the beam on each of the operating stations. The operators made use of the so-called light gun to tell the computer to associate a track with other keyed information, such as track number, identification, altitude, speed and armament. The operator placed the light gun over the display screen at the position of interest and pressed a trigger switch. When the screen was illuminated at that position, a signal was sent to the computer which said in effect that the deflection register contents identified the data item selected by the operator.

In order to reduce the load on the tracking programs, radar returns from fixed objects were filtered from the gap-filler data by a device called a video mapper. The mapper was a standard plan-position display for a single radar with a photocell viewing the whole display. Returns from fixed objects were covered with opaque material so that these returns did not activate the photocell and were rejected.

By the time the Cape Cod System was finished, it had about 30 operational stations with appropriate displays. The data required by an operator could not all be accommodated on the graphical situation display, so the Whirlwind group created an auxiliary display for text data associated with a particular track.

The Cape Cod System was used in exercises which included SAC bombers playing the role of hostiles, and the Air Defense Command and ARDC interceptors playing a friendly role. Before the experimental SAGE sector

which grew out of the Cape Cod System was finished, 5000 or so sorties had been flown against the system to test the system as well as its component parts.

WHIRLWIND II

It was clear to those who had participated in the Valley Committee and in the Charles Study that Whirlwind was more of a breadboard than a prototype of the computer which would be used in the Air Defense System. To turn the ideas and inventions developed on Whirlwind into a reproducible, maintainable operating device required the participation of an industrial contractor. The conceptual production computer became known as Whirlwind II.

The Whirlwind II group was set up in 1952 to deal with all design questions, including whether transistors were ready for large-scale employment (they were not) and whether the magnetic core memory was ready for exploitation as a system component (it was). The Whirlwind II group also spent much of its time in negotiation with Air Defense Command and ARDC Headquarters personnel in structuring the overall air defense system, including the definition of areas of control, cross-telling among sectors, need for weapons allocation, manning requirements, and air defense doctrine.

The most important goal established for Whirlwind II was that there should be only a few hours a year of unavailability of the operational system. The Whirlwind II team thought this was possible, extrapolating from the experience on the Cape Cod System. Most of the design choices faced by the Whirlwind II group involved the tradeoff among the number of tracks that could be processed, the number of interceptors that could be employed simultaneously, and the system availability criteria.

SELECTION OF A COMPUTER CONTRACTOR

The idea of engaging a manufacturer to help with the design engineering and manufacturing of the field computer was implicit in the nature of the R&D mission of Lincoln Laboratory. To achieve this end, a team was set up consisting of: Jay Forrester, Head of Lincoln Division 6, and of the Digital Computer Laboratory; Robert R. Everett, Associate Director of Division 6 and Associate Director of the Computer Laboratory; C. Robert Wieser, leader of the Cape Cod System Design; and Norman H. Taylor, Chief Engineer of the Division. They were responsible for finding the most appropriate computer manufacturer and designer to translate the progress made so far in the Cape Cod System into a design for the next generation air defense system. This system was to become known as the Lincoln Transition System, and in 1954 was renamed SAGE.

Early in 1952, this team made a survey of the possible candidates and

chose four for further evaluation. They were IBM, Remington Rand (two different divisions) and Raytheon. The team visited all three companies and reviewed their capabilities. They graded the companies on the basis of personnel, facilities and experience.

They looked at the technical contributions of the companies in terms of reliable tubes and other components, circuits, hardware, packaging, storage systems, and magnetic tape units. The companies were graded on their probable capability of bringing the Whirlwind II from development to production. This included their experience in setting up production of high quality electronics, their understanding of tests required, and the availability of their trained people. The team evaluated the production organization, the quality of assembly work, size of organization, similarity of the proposed work to the company's standard product, and present availability of production capacity. The team evaluated each company's service organization and training ability. Finally, the team considered the proximity to MIT and the train travel time to the various headquarters. Each of the four men on the team made his own assessment, using the weights decided upon before the trip. IBM received the highest score and was selected.

The IBM decision to accept the contract was made at the highest management levels. It involved evaluation of the risks versus the benefits. Some of the risks considered were technical feasibility, monetary risk, effect on commercial programs of losing people to this project, and potential liability for mishaps posed by the operation of a real time system. Offsetting advantages included direct involvement in technical advances plus an opportunity to respond to a national defense need.

THE IBM CONTRACT

From the point of view of the IBM people involved full time, Project High began in September of 1952 in anticipation of a study subcontract from the MIT Lincoln Laboratory. The subcontract was issued in October, covering a six month period. Office space was rented on the third floor of a necktie factory on High Street in Poughkeepsie, N. Y. The project got its name from this location. John Coombs was the first project manager. He had recently joined IBM from Engineering Research Associates.

During the next few months the expanding IBM group learned the current status of air defense studies. The expansion was done at the expense of other IBM development groups. The most important target for visits was the Boston area to study the Cape Cod System and to become acquainted with the overall design strategy of the Lincoln Labs people as well as their specific proposals for central processor design. A visit was made to a competing system at the University of Michigan. This was the Air Defense Integrated System (ADIS), which grew out of Project MIRO, a ground control system for the BOMARC ground-to-air missile.

In January of 1953 the system design began in earnest. IBM had bought the High Street building and had 26 people assigned. The Lincoln Whirlwind II

team organized itself along major subsystem lines. There was an arithmetic element section, a memory section, drum design section, and so forth. The IBM team organized itself in a similar pattern, and these counterpart groups began the work of trying to design the system on a joint basis. The Lincoln group, fresh from its experience of making Whirlwind I operate and designing the Cape Cod System, tended to view the IBM task as that of packaging Whirlwind devices so the system could be reproduced easily and quickly. On the other hand, the IBM people expected to participate in all levels of central computer system design and favored the technology they were familiar with.

The AN/FSQ-7 was designed by joint MIT-IBM committees that managed to merge the best elements of their members' diverse backgrounds to produce a result that advanced the state of the art in many directions. The committees presented their proposals at joint meetings that often involved 20 to 40 participants. Miraculously, these groups were able to arrive at a consensus and make progress. The MIT people had the final word on design specifications. However, most decisions really were based upon joint agreement.

During 1953 the design meetings involved a lot of traffic between Poughkeepsie and the Boston area. The roads had not been built for the needs of Project High, so driving was difficult. Some of the early meetings were held in Hartford, Conn., which was the half-way point between Poughkeepsie, N. Y. and Bedford, Mass. Another way of going from Poughkeepsie to Boston was to take an evening train to New York and a berth on the Midnight Owl to Boston. Small groups began chartering aircraft for a one hour direct flight, and on several occasions a large group would charter a DC-3. This helped to justify IBM's first corporate aircraft.

The first Hartford meeting was held January 20, 1953. John Coombs, the senior IBM man at the meeting, said that the purpose of the meeting was to allow the people working on the system, both at MIT and IBM, to exchange descriptions of what was being done. Jay Forrester, the first Lincoln speaker, went into some detail about the background of the program and his perception of the roles of the Lincoln and IBM people. He characterized the program as urgent, with a prototype system required by 1954. He referred to memorandum TM-20 which contained a description of what was then known as the transition system. He stated that none of the existing computers, including Whirlwind I, and the 701, and the others in existence were suitable. First, because of the nature of the problem, specialized peripherals would be required. Second, and probably more important, these existing machines had nothing like the reliability required for the job. Forrester suggested that IBM place a representative at the Cape Cod facilities. He gave a fairly complete description of the status of Whirlwind II thinking at MIT.

J. F. Jacobs of Lincoln presented the arguments for choosing vacuum tubes for the arithmetic and control units. It was too early for transistors and magnetic core circuits were too slow. H. D. Ross of IBM reported some tentative arithmetic element decisions, including the use of one's complement arithmetic and the use of flip-flops instead of the pulse regenerator used in the IBM 701. M. M. Astrahan of IBM described proposals

for logical design innovations. These involved index registers, dual arithmetic elements for simultaneous processing of X and Y coordinates of tracking data, and an interrupt scheme for operating in-out equipment simultaneously with program instruction execution.

Other Lincoln speakers were R. L. Best on basic circuits, W. N. Papian on magnetic core memory, J. H. McCusker on magnetic core production, and K. H. Olsen on the Memory Test Computer. Other IBM speakers were N. P. Edwards on non-memory magnetic core applications, E. H. Goldman on buffer storage and display, and J. A. Goetz on component reliability and standardization.

Lincoln's Norm Taylor discussed the schedule. He told the group that Lincoln had set an objective of having a prototype computer with its associated equipment installed and operating by January 1, 1955. Installation, testing and integration of the equipment in the air defense system had to be started on 7/1/54. The nine months preceding this, 10/1/53 to 7/1/54, would be required for procurement of materials and construction of the model. That left about nine months for engineering work in connection with the preparation of specifications, block diagram work, development of basic circuit units, special equipment design, and all the other things necessary to permit actual construction to begin. The schedule for this work was very tight. Taylor estimated that IBM would require about 235 development engineering professionals at the peak.

The meeting was concluded by T. A. Burke of IBM who described IBM's progress on the subcontract, which would end in three months. He was concerned that the follow-on Air Force prime contract be issued in time to avoid interruption of work.

A second joint meeting was held in Hartford on April 21. The first meeting had resulted in formation of a number of committees made up of IBM and MIT engineers who were to prepare design specifications. The second meeting consisted mostly of status reports from these committees.

In April of 1953 IBM received a prime contract for computer design specifications. On May 21, 1953, another Hartford meeting was held, this time to deal with packaging of Whirlwind II. Much of the meeting was spent on standardization of pluggable units. It was agreed that the mechanical design group should proceed with the design of a six tube pluggable unit, with backup designs for four- and nine-tube units. Another meeting on packaging was held June 1, 1953, at which a final decision was made to have both six- and nine-tube units. A breakdown of the central machine (arithmetic, control, and memory) into seven main frames was described.

Robert P. Crago joined Project High in June of 1953. He was to become manager of Engineering Design in July, 1954, and Manager of Project High in February, 1955.

PROJECT GRIND

The Hartford meetings acted as an information exchange, a catalyst for initiating action, an opportunity to identify overlooked aspects of the machine, and a forum in which people could interact on a personal level. By the time of the last Hartford meeting, a modus operandi had been established between the IBM and the MIT staffs. The central machine was pretty well agreed upon. It would have a single address order code in a 32-bit word. The memory would have a read/write cycle in the range of 5.5 to 7.5 microseconds for 8192 33-bit words, including a check bit. Data words required only 16 bits, so each retrieval involved two data words.

The central machine turned out to be the easy part of the job. In the rest of the system, decisions were not being made fast enough to meet the schedule. There was not enough time for detailed study of all the alternatives available, so choices had to be made primarily on the basis of the experience the individuals had with the subject area under consideration. To expedite this decision-making, it was agreed that a series of meetings would be held in which as many of the necessary decisions as possible would be made in a short period of time. These meetings were called Project Grind, since the participants were to grind away at each topic until a decision was reached. There were seven days of these meetings between June 24 and July 15, 1953. In order to identify the machine under design within IBM as well as MIT, the Whirlwind II name was dropped in favor of Air Force nomenclature, and the system was given an Air Force number, AN/FSQ-7. An AN/FSQ-7 planning group was identified, consisting of about 20 members drawn from both IBM and MIT. The procedure that was followed consisted of taking subsystems one at a time and forging whatever decisions could be made with the existing background and knowledge. Minutes of the Project Grind meetings were taken, to put on record some of the decisions made and some of the reasons for those decisions. Any problem could be brought into the open so that decisions could be made as soon as possible. It was also agreed that everyone should present even tentative plans for various parts of the system, so long as everyone knew that they were tentative.

The first Project Grind meeting, on June 24, 1953, was devoted to the radar inputs. Slowed-down video inputs, video mappers, and slowed-down video input registers were discussed. A general description of the input registers was agreed upon.

At the second meeting the subjects were marginal checking, power supplies and magnetic core memory. The third meeting dealt with magnetic drums. It was tentatively agreed that there would be six parallel fields of 34 bits each (two bits for status) per physical drum, with two heads per bit for input/output buffering, and probably five physical drums in the computer. The fourth meeting was concerned with output display systems. A two second display cycle was tentatively accepted. This meant that all display data streamed by the display consoles every two seconds and each console displayed the information requested by the operator. It was agreed that there would be 16 words available per displayed track, allowing for display of history of all tracks. The fifth meeting was concerned with cross-telling, output drums, output links for digital information, a

display maintenance console, and mechanical design.

At the sixth meeting the concern was standard circuits and the action of the Standards Committee. Four tube types were definitely approved, and it was decided that one-tenth microsecond pulses would be used wherever possible in the system. It was generally agreed that a project meeting should be held at least once every other week.

The seventh and last meeting, on July 15, covered mapper subcontracts, cross telling, review of the drums, paper tape machines, input counters, manual inputs and power supplies. It was generally agreed that paper tape would not be used in the FSQ-7.

DEVELOPMENT AND PRODUCTION

Project Grind resulted in fewer decisions than considered necessary to meet the schedule, but it had a remarkably good effect on the working relations of the people involved. It also demonstrated the need for some on-going method for reaching a consensus on high level specifications.

It was this need that eventually prompted Lincoln to set up a Systems Office under the direction of J. F. Jacobs, the purpose of which was then called design control. The Systems Office was in touch with IBM and MIT opinions. It was necessary for IBM and MIT to come to terms on the design of the FSQ-7, and it was also necessary that there exist a description, in specification terms, of what it was that the Air Force was buying. Thus the Systems Office took inputs from IBM, MIT, ADC (Air Defense Command), and Lincoln Project Office of the Air Force, and later inputs from the 4620th Air Defense Wing, and created a forum in which consensus about the main features of the design in all aspects of the system could be obtained. When this consensus was reached on the various parts of the system, a document would be prepared for the purpose of recommending to the Air Force that they approve or disapprove all or part of a proposed procurement of the pieces of the system.

In September of 1953 IBM received a contract for two single computer prototype systems, XD-1 and XD-2. XD-1 replaced Whirlwind in the Cape Cod System during 1955. The arithmetic, control and memory units were shipped in January to the Lincoln site in Lexington, Mass. Final testing was done there, along with integration of other frames which were shipped during the year. The modified system was renamed the Experimental SAGE System. The XD-2 was produced to support programming system development and to provide a hardware testbed in Foughkeepsie.

The broad outline of the SAGE network was delineated in 1954. The first serious plan visualized 46 computerized Direction Centers. It became evident to the Air Force that it would be desirable to automate the Air Defense Division headquarters. These headquarters, called Combat Centers, had the responsibility for directing the operations and allocating weapons on a large-scale basis, involving several Direction Centers. This called for a computer like the FSQ-7 with a specialized display system.

This system was named the FSQ-8. The locations for Q-7's and Q-8's were chosen and a delivery schedule was worked out calling for production of three systems the first year (1957) and ten to 12 in each of the subsequent four years. As the program continued, periodic revisions were made of the number of automated sectors and the installation schedule.

The first production contract was awarded to IBM in February of 1954. The first production system was accepted in its manufacturing test cell on June 30, 1956. The first system was declared operational at McGuire Air Force Base on July 1, 1958. To implement the deployment schedule, IBM built a manufacturing plant in Kingston, New York. IBM manufactured a total of 24 FSQ-7's and three FSQ-8's. These were deployed along the northern perimeter and the east and west coasts of the US.

INNOVATIONS

The SAGE system provides a demonstration of the kind of innovation that can be achieved when cost is secondary to performance. This kind of environment is difficult to create in a commercially-oriented company. However, SAGE provided the environment. Ambitious performance goals were met by the operational systems. Furthermore, as hardware costs dropped, most of the SAGE innovations became cost effective for the commercial market. The following items are highlights of some of these innovations.

- Core memory in a production machine. This is probably the single most important innovation in SAGE. The size and reliability required could not have been achieved by any other memory technology existing or proposed in 1953. The core memory used in SAGE evolved directly from the pioneering work of Jay Forrester and the MIT groups that developed the feasibility model and built the Memory Test Computer (MTC). "By May 1953, the MTC was demonstrating the swift, highly reliable operation of arrays of cores 32-by-32, stacked 16 high" [8]. The original system design called for 8192 words of 33 bits, including a check bit, arranged in two banks of 33 planes. Each plane was a 64x64 matrix. When the requirements of the application program became apparent, a 256x256 unit (65,536 words) was designed to replace one of the smaller banks. In cooperation with the MIT group, IBM developed the methods of manufacturing and testing uniform reliable and inexpensive core memory in production quantities. This involved an automatic core tester and a core plane stringing process which used hypodermic needles to guide the fine wires through the tiny cores.
- Active/Standby Duplex System. The AN/FSQ-7 was the first computer system to use two computers in active/standby roles for reliability. Previous dual computer systems had both computers doing the same thing and compared output. In SAGE, the standby computer could run test programs or other work while the active computer ran the air defense programs. The active computer maintained situation status information on an intercommunication drum accessible by both computers [3], [10].

The original concept was three computers located at different sites within a geographical area called a sector. The radar inputs were to be connected to two of the three with sufficient displays so that any two of the three could run the system at full capacity. This mode was rejected because of the high communication costs and high costs for replicated personnel support facilities.

The duplex decision was not made until November 1953. Since it involved design changes in the I/O components, a separate group was formed to do the redesign without affecting the schedule for construction and test of the XD-1 and XD-2 prototypes. The design philosophy was to duplicate every unit that could shut down the whole system. Thus the central computer and input drums were duplicated, but display consoles and modems were not [3]. Great care was taken to ensure that the switchover facilities did not introduce single failure modes affecting both of the duplexed systems.

- Digital communication over standard phone lines. The transmission of digital data over voice-grade phone lines at 1300 bits per second was pioneered by the Lincoln people. Jack Harrington's group of Division 2 designed the first modems to convert digital data to and from analogue waveforms that could be accommodated by voice-band channels. The channels required special conditioning to minimize noise pickup and eliminate unequal phase shifts across the frequency spectrum. The phase shifts were not noticeable in voice transmission but distorted the data waveforms.

- Time sharing. Time sharing a computer for real-time tracking of hundreds of airplanes, real-time control of weapons, and interaction with human controllers was a bold concept. It required invention of programming techniques to ensure timely sequencing through all the tasks [3]. Programs and data tables were paged in from drums and the tables only were rewritten. Data input/output and display data were fully buffered by drums.

- I/O control with memory cycle stealing. SAGE marked the introduction of the I/O break, also called memory cycle stealing. This forerunner of modern channels allowed computation to continue during I/O operations, interrupted only for the core-memory cycle required to transfer a word between the core memory and the I/O device [3]. It involved a register to count the number of words transferred and a memory address register, incremented for each word transferred, to specify the location of the next word [2].

- Associative input system with drum buffer. The input buffer drums contained radar data from several sites intermixed. Each data item was tagged with the identity of its source radar. The CPU could request all the data from a particular radar. This constituted an associative memory access.

- Branch and Index instruction. The AN/FSQ-7 index registers were an adaptation to a parallel machine organization of the Williams B-tube [11]. The Branch and Index instruction allowed a single instruction to decrement an index register, test for the end of a loop, and branch

back to the beginning of the loop [2].

- Computer control of marginal checking. Marginal checking by varying supply voltages was proven effective for vacuum tube circuits by the Whirlwind experience. The AN/FSQ-7 extended the capability by allowing program control of the voltage excursion magnitude and its point of application [1].
- Display, Light Gun and Keyboard input in a production machine. The Cape Cod system demonstrated the functions needed in a cathode ray tube display console, including the use of light guns. The AN/FSQ-7 display system constituted the first use of such consoles in a production computer system. The graphical situation displays used the Convair 19-inch Characteron tubes in which the electron beam was passed through a mask in order to shape the beam into the form of one of 64 characters. The shaped beam was then deflected to the desired position on the screen. There was also a textual display which used the five-inch Hughes Typotron. The Typotron also used a character mask but had a storage screen instead of the standard phosphor. IBM designed the display consoles but subcontracted production to Hazeltine.
- Circuit standards. A central circuit design group was responsible for design or approval of all CPU circuits. The group followed a set of design standards based on component tolerances and compatibility with marginal checking [6].
- Component specifications and vendor control. Special contracts were made with manufacturers of vacuum tubes, capacitors, diodes and resistors to assure the uniformity and reliability of the products they delivered. IBM required these vendors to institute strict controls over the design, manufacture and testing of the components and actually monitored the manufacturing and testing processes at these vendors' plants [4].
- Circuit packaging. In the pluggable units, all components except vacuum tubes were mounted on etched circuit boards. IBM's Manufacturing Engineering Department worked with General Mills to develop the Autofab machine, which assembled and soldered the circuit boards. These automatic soldering techniques greatly increased the reliability of the circuit boards, as did the development of double-sided boards with plated-through holes.

ILLUSTRATIONS

Figure 1 shows the components of a SAGE sector external to the Direction Center. Figure 2 is an external view of a Direction Center. Figure 3 shows operators at the situation display consoles. Figure 4 shows a light gun in use. Figure 5 shows the computer control room, including the operating consoles at the left and right and the duplex switching console at the far end. Figures 6 and 7 show the front and back, respectively, of a

typical computer frame. Figure 8 shows a 4096-word magnetic core memory unit. Figure 9 shows a magnetic drum frame.

POSTSCRIPT

An AN/FSQ-7 system weighs 250 tons and has a 3000 kw power supply. The first Direction Center began operating in 1958, and there are seven still in operation. Performance data was compiled on these seven for the 24-month period from March 1978 to February 1980. Each system uses 49,000 vacuum tubes. The tubes have an MTF of 50,000 to 100,000 hours. The average percentage of time that both machines of a system were down for maintenance was 0.043%, or 3.77 hours per year. The average percentage of time both machines were down for all causes, including air-conditioning and other situations not attributable to the computers, was 0.272%, or 24 hours per year.

REFERENCES

1. Astrahan, M.M., and Walters, L.R. "Reliability of an Air Defense Computing System: Marginal Checking and Maintenance Programming" IRE Transactions on Electronic Computers EC-5,4 (Dec. 1956) 233-237.
2. Astrahan, M.M., Housman, B., Jacobs, J.F., Mayer, R.P., and Thomas, W.H. "Logical Design of the Digital Computer for the SAGE System" IBM Journal of Research and Development, 1,1 (Jan. 1957) 76-83.
3. Everett, R.R., Zraket, C.A., and Benington, H.D. "SAGE--A Data Processing System for Air Defense" Proc. Eastern Joint Computer Conference (Dec. 1957) 148-155.
4. Heath, H.F. Jr. "Reliability of an Air Defense Computing System: Component Development" IRE Transactions on Electronic Computers EC-5,4 (Dec. 1956) 224-226.
5. "MITRE The First Twenty Years" The MITRE Corporation, Bedford, Mass. (1979).
6. Nienburg, R.E. "Reliability of an Air Defense Computing System: Circuit Design" IRE Transactions on Electronic Computers EC-5,4 (Dec. 1956) 227-233.
7. Oglatree, W.A., Taylor, H.W., Veitch, E.W., and Wyles, J. "AN/FST-2 Radar Processing Equipment for SAGE" Proc. Eastern Joint Computer Conference (Dec. 1957) 156-160.
8. Redmond, K.C., and Smith, T.M. "Lessons from Project Whirlwind" IEEE Spectrum 14,10 (Oct. 1977) 50-59.
9. Redmond, K.C., and Smith, T.M. "Project Whirlwind, The History of a Pioneer Computer" Digital Press, Bedford, Mass. (1980).
10. Vance, P.R., Dooley, L.G., and Diss, C.E. "Operation of the SAGE Duplex Computers" Proc. Eastern Joint Computer Conference (Dec. 1957) 160-172.
11. Williams, F.C. and Kilburn, T. "The University of Manchester Computing Machine" Review of Electronic Digital Computers (Feb. 1952) 57-61 (Joint AIEE-IRE Computer Conference, Philadelphia, Dec. 10-12, 1951)

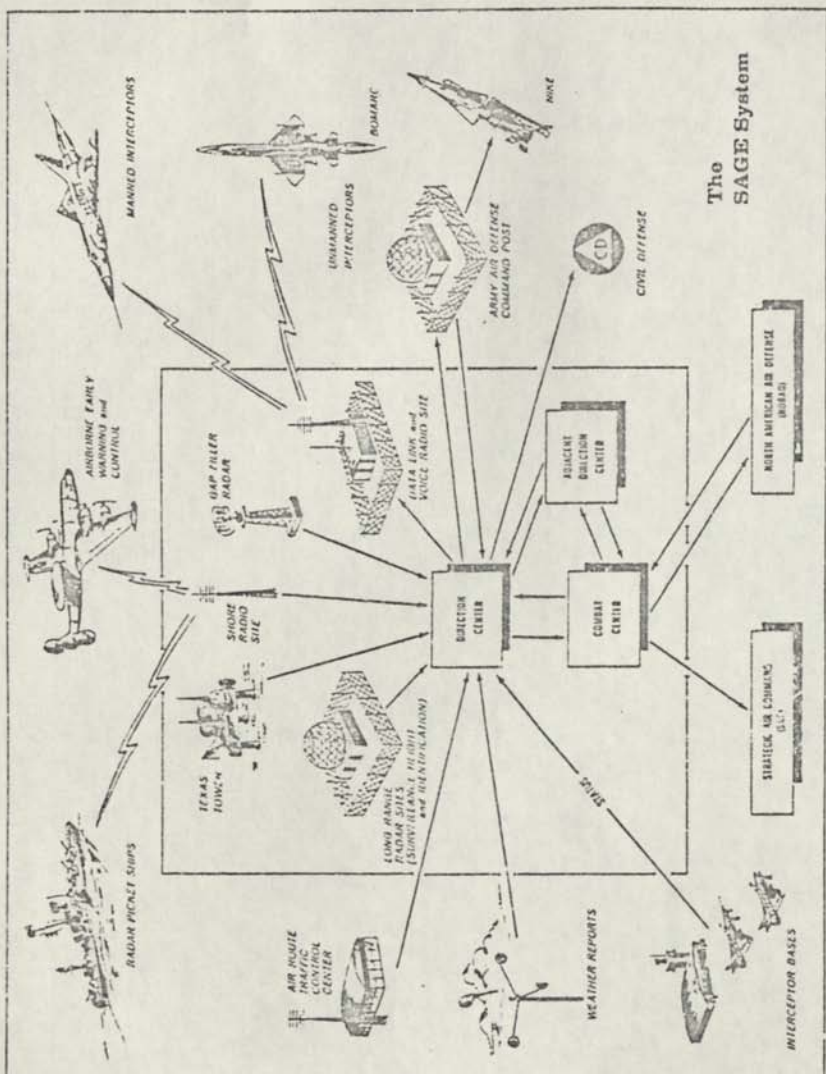


Figure 1




Figure 2



Figure 3





Figure 4




Figure 5




Figure 6




Figure 7

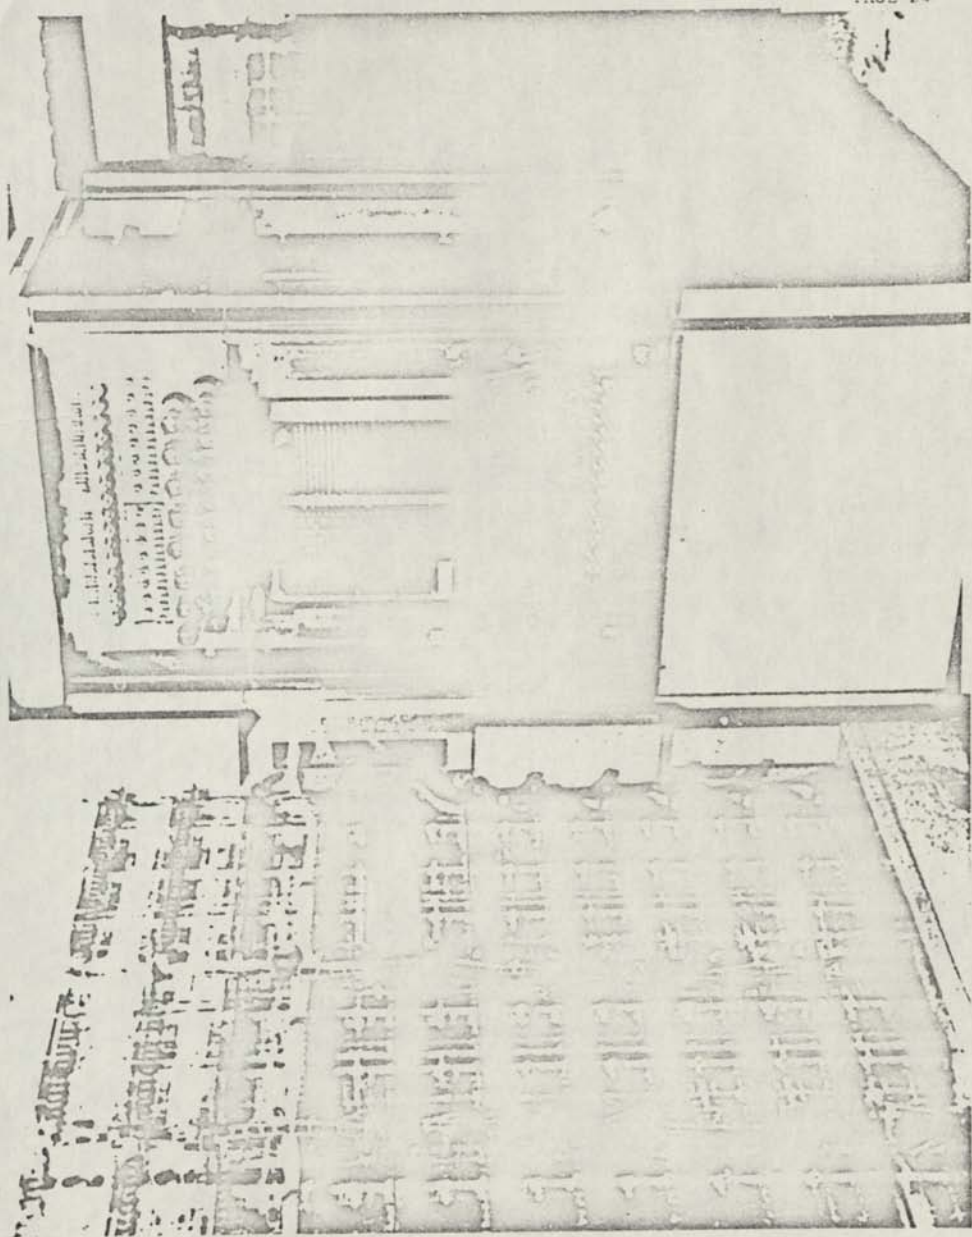


Figure 8

