Whirlwind's Genesis and Descendants

"Worldwind's Genesis and Descendants" was the theme of a symposium held at The Computer Museum October 18, 1987. This was part of a weekend reunion of the Whirlwind group organized by David Israel. The symposium was recorded at the Museum and transcribed by Judy Clapp of the MITRE Corporation. Responsibility for the accuracy of the following adaptations of the talks belongs to The Computer Museum.

Whirlwind's Success

Jay Forrester

Jay Forrester is Germeshausen Professor of Management and Director of the Systems Dynamics Group at MIT. He was the leader of the Whirlwind group at MIT from the late forties until 1956.

Jay Forrester, T. K. Finletter, and F. Wheeler Loomis visit the Whirlwind in November, 1951.

Why did Whirlwind succeed? Why did more technical innovations out of Whirlwind persist into the present time than from any other of the early computers? The reason revolves around several things: the vision of the future direction of computing, a dedication to excellence, and the organizational environment.

Project Whirlwind's Future Vision
The vision in Whirlwind reached well beyond the uses of computation and hand-calculating machines at that time. Our work quickly became identified with the field of real-time control and reliability.

The dedication to real-time control started well before Whirlwind first operated. In October 1947, when we were still determining the logical structure of the machine, two reports were written in the MIT Computer Laboratory suggesting that the Navy could use digital computers as Combat Information Centers for co-ordinating an anti-submarine task force. This meant coordinating the air, the surface, and the subsurface pictures to get an understanding of the totality of what was going on.

Building Reliable Systems
Reliability was important because you can't go back and do things over again in military applications. In 1948, before Whirlwind operated, Karl Compton, then President of MIT and also Chairman of the Research and Development Board, asked that we prepare a memorandum for him on the future use of computers in the military. Bob Everett, Hugh Boyd, Harris Fahnestock and I took two or three weeks to answer that question. The report culminated in a chart listing vertically about twelve wide-ranging areas of computer use in the military, such as logistics, scientific computation, air defense and anti-ballistic missile control. On the other axis were 15 years from 1948 to 1963.

That report is quite an interesting document in historical perspective. At each intersection in each square in the table, we estimated the condition of the field at that time, how much money would be spent yearly in research, engineering and production, and what the condition of the field would be relative to those end uses 15 years into the future. These estimates were made when no high speed general purpose computer had yet functioned.

The estimates are percentage-wise as good as and maybe better than most estimates made today for the time and cost of the next computer to be put into production. This was because we paid a great deal of attention to the political as well as the technological side. The cost estimates were arrived at by subdividing tasks to no more than 30 people working a calendar quarter and by deciding all the things that would have to be done. It was not necessarily correct in detail but it was a logically complete scenario including how long it would take for people to believe the results of the previous year, and how long it would take to get funding for the next step. The chart showed a total of $2 billion to be spent in research and development alone over the 15-year period. We went into a Navy conference with this. They thought the agenda involved whether we could have the next $100,000. There was a communication gap in that meeting.

Dedication to excellence
Many people in the Whirlwind group had had the World War II experience of going from theory through research to production design, then to manufacturing and into the battlefield, fixing their own mistakes at every stage. They understood how the decisions at the research stage really affect what happens later.

In my own early background, I had already started down that road, having grown up on a cattle ranch where you learned that if you did a sloppy job of fixing a tractor or a well, you would suffer the consequences very soon, have to do it over, and do it right. Part of the manifestation of that viewpoint showed up, of course, in our improving vacuum tubes. Until the 1950s, vacuum tubes primarily had been used for radios. Radio engineers were not concerned that the life of a vacuum tube was about 500 hours. But computer engineers, considering the use of many thousands of vacuum tubes, easily estimated that with such a short life, the machine would run no more than a few minutes between failures. One of the achievements of our group was determining the cause of failure of vacuum tubes. It turned out to be one thing. After removing that cause in the design, the life of vacuum tubes was increased, in one design step, from 500 hours to 100,000 hours or longer.

Excellence also meant thorough testing of components. We built a five-digit multiplier for the simple purpose of finding out whether an electronic device running continuously would be trouble-free or not. There was uncertainty about things that people now thoroughly understand.

One important issue was our uncertainty about thermal noise. We didn't know if random spikes of thermally generated noise were big enough to trigger our robust computing circuits. We wondered whether thermal noise would intrude itself often enough to be devastating to accurate computation. To test for this, the five-digit multiplier was run continuously. Every multiplication was checked against a reference number. Sure enough, it didn't compute reliably all the time. It had a great tendency to make mistakes at 3 a.m. This was traced to the janitor in the building next door, who would start the freight elevator at about that time, upsetting the power circuits enough to produce a computation error. As a result, a rotating motor generator with enough inertia to carry through that kind of transient noise was installed on both Whirlwind and the SAGE Air Defense machines. It was an expensive solution but a very effective one.

A lot of time was spent writing test programs to find out the source of a failed component. Occasionally, a visitor was asked to go any place in the computer racks, pull out a vacuum tube and bring it back to the control desk. When he got back, the location of the empty socket would have been typed out by the machine itself. Finding solid, existing, reliable errors, like a tube pulled out of its socket, was not nearly good enough.

Other means of determining reliability were also essential, which we discovered in various ways. I remember one Saturday, during one of many annual reviews, our inquisitor asked, "What are you going to do about the electronic components that are drifting gradually and are on the edge of causing mistakes? Any little random fluctuation in power, or streetcars going by, will cause circuits to sometimes work and sometimes not." This was a very important and powerful question that, frankly, we had done nothing about. It was such a pointed question and obviously such an important one that I felt an immediate answer was essential. I said to him, "Well, we could lower the voltage on a tube and convert it from a marginal to a permanent failure and then it would be easy to find." He thought it was a good solution and so did we, so the next Monday we started designing it into the computer. The marginal checking system in Whirlwind carried over into the SAGE Air Defense system, adding another factor of ten to the reliability.

Many of you may not know the statistics on the SAGE system's reliability. There were 30 or more SAGE Centers. Each building was about 160 feet square, four stories high, with upwards of 60,000 vac-

uum tubes in it. The question is: what percentage of the time do you think such a center would operate reliably? The answers I get from an audience today tend to run from 15% to 60 or 70%. They're really quite overwhelmed when they're told the historical statistics on the SAGE Air Defense system. It was installed in the late 1950s and operated for 25 years, until 1983. According to the data that Bob Everett was able to find, the uptime was 99.8%, which is really quite remarkable. In fact, you will have trouble finding anything equal to that, even when it has been designed with more modern components.

The attitude about the SAGE performance was that it must work reliably. To achieve high reliability, one must be a devout believer in Murphy's Laws — that if anything can go wrong it will. Every possible failure must be identified and forestalled. This attitude is the difference between something that is strikingly successful and disaster. In almost any major disaster, whether a technological or a social one, an ample number of people knew that it was likely to happen and knew in advance why it was going to happen. The information was there, and either they did not take any action, or they tried, and in the social circumstances of their environment, were not able to get any results. A warning is almost always present ahead of the trouble and the problem comes in getting any kind of action or acceptance of the threat.

The Organizational Environment
Another part of the success of the Whirlwind group came from the organizational environment within which we were operating. MIT in those days was a free enterprise society in which someone who had a vision and could raise the money for it could do what he thought was important.

The Leaders
Within our immediate environment, two

people conspicuously stand out as having made it possible for us to operate the way we did. One was Nathaniel (Nat) Sage, Director of the Division of Industrial Cooperation, under which outside funding came into MIT and the other was Gordon S. Brown. In addition, there were two promoters, in the best sense of that word, people who shared the vision and who spent their time building up the outside constituency to support the work. These were Perry Crawford and George Valley.

Sage, a civil engineer by training, was the son of an Army officer and grew up in Army camps around the world. Somewhere in that experience, he developed into a very good and self-confident judge of people. There were people at MIT that he trusted implicitly, and there were others that he wouldn't trust any farther than he could see them. Sage trusted Gordon Brown, Stark Draper, of the Draper Laboratory, and I think I can claim that he trusted me. He had confidence in us, lent great support to us, and would do rather remarkable things for us. I remember when someone chartered an airplane to come back from somewhere because it was a sensible thing to do to get home for the weekend. That caused an explosion in the Military Contracting Office where they thought this was not an appropriate use of funds. The contracting officer went to Nat Sage as the senior person. Sage would listen to them, nod, sympathize with them and say, "That really is too bad." Then he would put the whole thing in his desk drawer. He would never even tell us that the question had been raised, because he believed it probably was a proper thing to do.

Gordon Brown, my mentor at MIT, and director of the Servomechanisms Laboratory under which the Computer Laboratory operated, was a person who threw a great deal of responsibility onto young staff members, even as research assistants in the Electrical Engineering De-

Dr. Gordon S. Brown, who was Director of the Servomechanisms Laboratory, helped create an atmosphere in which the Whirlwind project could succeed.
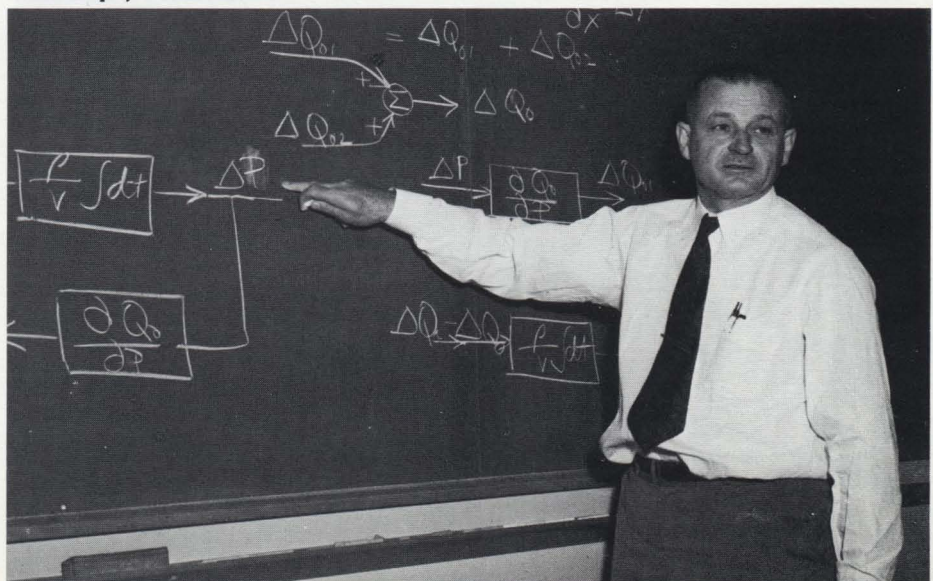
partment. He provided an environment in which people developed very rapidly, and in which they could attach themselves to some important and overriding goal. To him, goes much of the credit for making the environment where the Whirlwind computer project could flourish.

In 1939, Perry Crawford did his MIT Master's thesis on digital computation, which meant developing a ten-stage ring counter to compute with decimal numbers, but never carrying it beyond some individual computing circuits. He is a philosophical, looking-into-the-future type of person. By the time we made contact with him, he was in the Special Devices Center of the Navy in Port Washington, Long Island.

Perry Crawford is the person who first called my attention to the possibility of digital computation. We were standing on the front steps of 77 Massachusetts Avenue one afternoon when we were still working on analog computers in the Servomechanisms Lab. He began to tell me about the work on the Harvard Mark I computer, and about the ENIAC computer which was then under construction. He was a very uninhibited, unbureaucratic type and would circulate freely right up to the Naval Chief of Operations even though he was a civilian far, far down in the organization. He moved through the Navy selling the idea that digital computers had a future as Combat Information Centers. He had several computer projects under his direction that he raised money for. He is also the person who gave Whirlwind and other projects their names. All of them were named after air movements: Hurricane, Zephyr, Typhoon and Whirlwind.

The other promoter to whom we owe a great deal is George Valley, a professor of physics. He was on a committee of the Air Force looking into air defense. In the later stages of our work that led into

From left to right: Jay Forrester, Norman H. Taylor, John A. O'Brien, Charles L. Corderman, and Norman H. Daggert inspect the open, high voltage Arithmetic and Electrostatic Storage Racks characteristic of computer equipment in the early 1950s.

Lincoln Laboratory, he was the person who would call up generals in the middle of the night, tell them what they should do, and ask for support. He did all those things you read exposés about in books on the politics of technology, but which are necessary to keep the program coordination running smoothly.

The Organization
Sometimes you have people in an organization, each of them with an IQ of 130, and come out with an organization whose IQ is 70. What you get is the least common denominator rather than the best of the participants. I'm not sure how one creates the opposite environment, but there is great power in a tightly knit organization that has the capability of using the strengths of each person and compensating for the weaknesses of each.

Every person has strengths and weaknesses. You need a team in which there are such things as a vision of the future, a sensitivity to political matters, the capability of developing people, technical competence, the courage to transcend adversity, salesmanship, integrity, and putting long-range goals ahead of the short term. We had those characteristics well represented, scattered throughout our group. No person had all of them. For every person there would be, perhaps, a glaring hole in one of those dimensions. Yet, it was a group that understood each other well enough to use people in situations where their strengths prevailed rather than their weaknesses. Out of that came an organization that was able to be much more effective than most of those we see around us in technology and in most corporations at the present time. It is still an unsolved challenge to understand how that sort of spirit and unity can be created.

The Hostile World
Another thing that helped us, but that we resented, was the hostility towards innovation. There was little outside understanding of our subject, the objectives, or the methods for building pioneering computers. Funds were almost always inadequate. Reviews and investigations required us to defend our position and to face the weaknesses that other people were pointing out. We benefited from the distractions caused by the periodic reviews in which everything was questioned. Why were we using so much money? Why were we running late? Why were we designing the machine the way we were?

The matter of cost was one of the things that the outside world understood least. Whirlwind was being judged in the context of mathematical research, in which the salary of a professor and a research assistant was the standard by which projects were measured. We were spending way beyond that level, and

The Whirlwind console room in 1951 with the marginal checking and toggle-switch test control panels on the left. Stephen Dodd, sitting at an input device, is being watched by Jay Forrester and Bob Everett. Ramona Ferenz is seated at the prototype display to monitor the Cape Cod system, the prototype for SAGE.
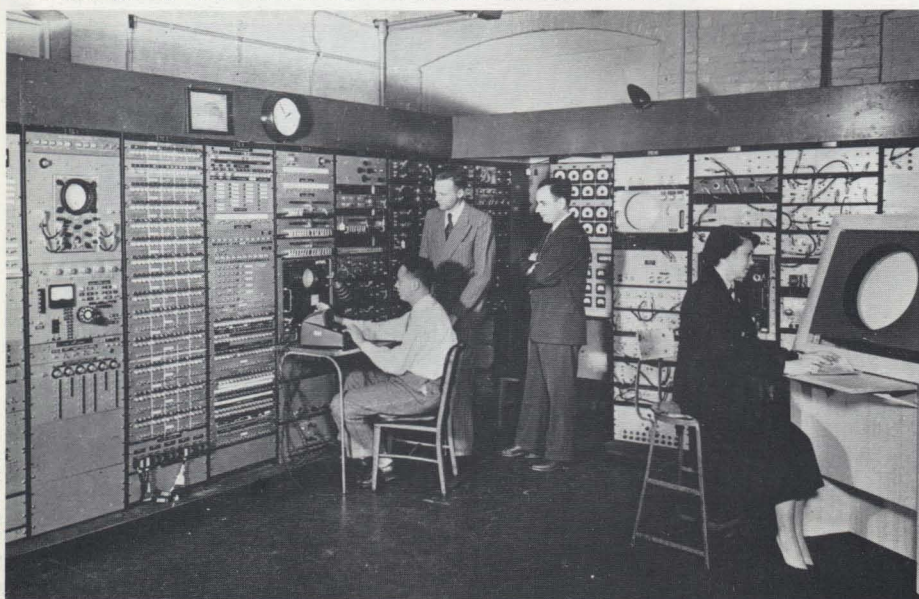
were seen as running a "gold-plated operation." Although the gold plating was occasionally excessive, in retrospect, I think there was reason for it.

An organization can't run with two contradictory standards. If you're going to have high performance and high quality in the things that matter, it is very difficult to have low quality and low performance in the things that, perhaps, don't matter. For example, at an early demonstration for important people, we didn't want them sticking their fingers into the high voltage in all those racks of Whirlwind. I asked somebody to get rope to put along the aisles so visitors wouldn't walk among the racks of vacuum tubes. A nice-looking white nylon rope was procured and installed. During the demonstration, I saw some of our critics fingering this beautiful rope and looking at one another knowingly as if to say, "That's what you would expect here." It may not have cost any more than hemp rope, but it reinforced that impression of an extravagant operation. Another example was the Cape Cod display scopes built into plywood cabinets faced with mahogany. Although our cabinetmaker made these quite inexpensively, people looking at those mahogany cabinets, were reinforced in thinking we were extravagant. Eventually we solved this problem by spending additional money and painting the cabinets gray.

Whirlwind's Technology
Making the decision to build Whirlwind I with a 16 binary digit register length was tremendously hard for us. The mathematicians were up in arms. They thought it was too short to be of any possible use. We defended it at that time on the basis that it was a demonstration of feasibility and we would build a 32 or a 36 bit computer when the right time came. Many of today's desktop computers are still 16 bits and only now moving to 32 bits. Selecting 16 bits was not a useless register length for computing, only a serious short term political problem.

The objectives of a computer at that time dominated the kind of high-speed internal memory to be chosen. Since Whirlwind was for demonstrating a very high speed computation for real-time applications, we chose electrostatic storage tubes rather than any of the more reliable kinds of serial memories. Each electrostatic storage tube with 1024 binary digits cost us about $1000 and had a one month lifetime. That meant that the upkeep on a storage tube, just its replacement, cost about $1 per binary digit per month. If you were to spend that on your two-megabyte personal computer, it would cost you $24 million per year just to maintain computer storage. The improvement has been perhaps a million-fold since that time in cost. That's about a factor of two every two years in the intervening 40 years. The high cost of storage tubes was the major incentive for inventing and perfecting coincident-current, random-access magnetic memory.

The economy necessary in programming was quite remarkable by today's standards. We demonstrated a military combat information center with one real bomber, one real fighter, and a radar set to generate data, with the computer receiving radar data by telephone line, analyzing it, throwing away the noise, averaging and smoothing and predicting the track, doing the same for the fighter, computing the intercept heading for the fighter, and then transmitting instructions to the autopilot automatically. If we today asked a programmer how much computer memory would be necessary for such a program, the programmer would probably guess a million bytes, minimum. The task was done on Whirlwind with 650 bytes of memory, not megabytes, just plain bytes. It was a time when the costs favored cutting programs to the minimum and using, if necessary, a lot of time, a lot of manpower, to reduce the programs.

Contributions of Whirlwind
In spite of the sense of extravagant expenditure, the entire Whirlwind project totaled about $4,500,000. That doesn't seem like much in today's computer world. Out of that came the first parallel, high-speed, clock-driven computer, magnetic core memory, cathode ray tube displays driven by a computer, an interactive light gun connecting a person to the computer, and many other innovations that are still important today.

We thought we had a good view of the future and we did for the succeeding 15 years, but I must say that our view of the future did falter if you were to extend it beyond that time. I gave a talk in the mid-1950s to a computer convention in which I pointed out that the cost of computation had been falling by a factor of

two every two years from 1940 to 1956. I said, "Of course that can't go on for very much longer." But, of course it did, and is still going on.

Becoming a User
After 1956, I went more into the use of computers, using the ideas of feedback systems that Gordon Brown had originally pioneered and applying the methodologies and concepts to understanding the behavior of social systems. My present work is focused on the way in which the policies of a corporation produce its successes and failures and the way in which the policies embedded in the private and governmental sectors produce the behavior of the national economy.

My present work is focused on understanding the so-called economic long wave, the great rise and fall of economic activity with peaks every 45 to 60 years. This behavior has produced the great depressions of the 1830s, the 1890s, and the 1930s. We believe that the present economic cross-currents are the beginnings of another such major downturn. Working on behavior of social and economic systems is now especially timely. Just as the frontier of physical science opened up in the 1800s, the frontier of understanding our social systems now lies immediately ahead.

The Whirlwind project had shown that a reliable real-time computer could be constructed and that aircraft could be tracked and intercepted. Robert Everett is shown here on the Control Force Demonstrator in 1947.