May 21, 1984

The History of the Advanced Scientific Computer

Harvey G. Cragon

In 1965, the 870 project had been concluded in Dallas. The project had been turned over to Len Donohoe's activity in Houston for upgrade engineering and production. A decision was made to start a study of a new computer which would significantly increase the computational capability of GSI for seismic processing. Milo Backus had become convinced that an interactive environment was needed to replace the batch environment then in common use. Milo was requesting that we now call it a "personal computer". The capability he wanted is just now coming on line, some 20 years later.

In January 1966 a task force was established to study the "Next Generation Computer Organization & Construction" as well as "The Role of Integrated Electronics and New Memory Systems". This task force was headed by Harvey Cragon and consisted of Len Donohoe, Dick Abraham and Jack Kilby. A similar task force under Milo Backus was to study computational requirements and growth potential for the 827 and 870 computers. Willis Adcock was in overall control of these two task forces.

There was activity in CRL on associative memories at that time. Willis arranged contacts with Cuthbert Hurd and Sid Fernbach to discuss this technology. These contacts lasted throughout most of the life of the ASC program.

At the same time we started this work, Dan Slotnick, at The University of Illinois, was starting the early work on what was to be known as the Illiac IV. I think that Willis Adcock was responsible for starting the contact we had with Slotnick. He had an $8M contract from ARPA to develop this machine. Dan visited TI on 1 April 1966 to look at our work.

On 21 April, we had made the decision to proceed with a serious study on a high speed computer. We named our program the TI Parallel Processor Study or TIPPS. We received $250K from Mark Shepherd to do this work. Contributors on this study were Harvey Cragon, Joe Watson, Ed Husband, Bill Kastner, Tom Cooper and Tony Galindo. Several reviews were held with Slotnick in May, June and July.

During the study, we gave very serious consideration to the architecture of the Illiac IV. This architecture consisted of an array of processors, each having its own memory but each sharing a common control unit. This architecture seemed ideal for processing well ordered arrays of data and should be low cost in production because of the replication of identical units. Nevertheless, the deeper we got into the evaluation, the more convinced we became that there would be severe programming problems and that as a "general purpose" super computer, the performance would not be impressive.

The other architecture we considered was a pipelined vector machine. We were influenced by the work of Seymour Cray, Leonard W. Cotton, D. N. Senzig, R. V. Smith, and Gene Amdahl's work on the IBM 360/95. By the end of July it had finally dawned on us that the Fortran "Do Loop" was an automatic invocation of a vector instruction, and our decision was made. We would adopt the pipelined vector architecture because we could build a vectorizing Fortran compiler. I recall that I was very fearful of our chances of convincing our management that we were correct and Dan Slotnick was wrong, as he was receiving national recognition for his work. We had been given a small, $50,000 contract by Slotnick to study these architectures. When we made our decision to go our own way, Willis Adcock and I went to see Dan to tell him the news. I recall that we saw Dan in a motel room on Sunday night in Champaign, having made the decision to see him the night before. Dan understood our position and encouraged us to continue our work. He encouraged us to try to get support from Sid Fernbach at Lawrence Livermore Laboratories (LLL). Sid was very interested in vector computers and was planning to get AEC funds for an RFQ in the fall of 1966. We did and this marked the start of a very long and unsuccessful attempt to sell a machine to the AEC.

Another decision had been made by early August on the architecture of the ASC: the ASC would be a multiprocessor. We had observed the significant performance improvements achieved on the TIAC 827 AND TIAC 870 when vector boxes were added by the Houston group. Further, Seymour Cray had effectively used peripheral processors on the CDC 6600, and we felt that this arrangement represented a good division of the work to be done. That is, operating system and I/O function in the PPU's and the number crunching in the central processor. A serious investigation was made into the possibility of using a low end model of the IBM 360 line as the peripheral processor. We rejected this idea due to the low I/O capability it would provide. Seismic processing required high I/O bandwidth and we could not get it from the 360.

During October and November, the TIPPS effort continued and a formal review was held on 5 December for Mark Shepherd, Fred Bucy, Cecil Dotson, Bob Dunlap and Mark Smith. The system specifications and development program for a seismic computer were to be reviewed. It was anticipated that the prototype would be completed in 1969 and the first production unit would be available in 1970.

The year 1967 was one of heavy activity in fleshing out the architecture. The memory system would consist of four thin film memories interleaved by means of a crossbar switch. This memory technology was the fastest available at that time and was also planned for the Illiac IV. The decision was made to implement the arithmetic pipeline as a reconfigurable pipe rather than use the Cray technique for a multiplicity of dedicated pipes. The reason for this decision is not clear to me at this time, but was probably made in anticipation of a hardware savings. We also departed from Cray's approach to the Peripheral Processor. Cray had used ten processors each having equal time access to the ALU. We chose to have eight processors having a dynamically assigned access to the ALU. We decided to have a deterministic system with no interrupts. All control function would be provided by a polling system in the PPU's. The system was provided with only two interrupts. These were for power going down and an arithmetic exception. The latter was maskable. Power going down is the only event which must be serviced; thus, it was given an interrupt.

Late in 1967, Joe Watson and Ed Husband developed the idea of the Communications Register Unit (Appendix A). They had spent a great amount of time on the TIAC 870 working the problem of the special I/O device control boxes. They wanted to have a system for the ASC of programmable I/O control. Thus, the ASC PPU's were viewed as micro machines which would serve several functions, depending upon the code which each executed. In this context, the CRU was developed with the bit pusher instructions. Recent events with various PC's tell us how important this idea was. It was also in the PPU design that we departed from a Cray approach. Cray had a private memory for each PPU. To start a task in a PPU, the task had to be loaded from main memory. We felt that a better way to operate would be for each PPU to have access to main memory without a local memory. Memory bandwidth would have to be high enough to support this activity.

Various reviews were held during the year 1967. During February, March and April, the desire for an outside review of our program became evident. We therefore contracted with Cuthbert Hurd to conduct such a review. Hurd's company CUC gave us their report in May. This contact lead later to CUC's taking on a major share of the early software development. In May, we estimated that the 1967 effort would cost between $600K and $1M.

A major debate between the project people and GSI concerned the word length of the computer. The project wanted a 32 bit word: GSI wanted 24 bits. This issue was resolved in favor of 32 bits in April.

Emory Garth had joined TI in late 1966. He and Gene Lee developed the major concepts of the ASC circuits and packaging during 1967 (Appendix B). A significant activity concerned bringing the design automation people into the effort to design the very complex printed wiring boards. This effort in DAD was quite successful and made a major contribution to the technical success of the ASC (Appendix L).

By the end of 1967, most of the major design decisions had been made and the program started to ramp up for a year of heavy detailed design in 1968. CUC personnel joined the program in Dallas during the second half of 1967 to start the software development program.

During May of 1968 we gave presentations on the ASC to Dr. Roger Lazarus and his staff at Los Alamos. This was a continuation of our effort to sell a machine into the AEC.

Thermal mockups of the ASC were built and tested during the early months of 1968. These tests indicated that we would be able to cool the system, and detailed design was started. The construction of the Peripheral Processor began in late 1968. As this unit was needed to run the operating system and control the peripherals, it needed to be constructed and checked out first. A move to the top floor of the South Building occurred in 1969, and it was here that the checkout began of the PPU.

By summer, the design of the Memory Control Unit and the Central Processor Unit began. Our first plans were that the MCU would be asynchronous between the memory modules and the processors. This was a design decision we lived to regret, as we later had to redesign for synchronous operation. Also during the summer, we began to get the first feelings that a one pipe machine would not be competitive in the market against the CDC 7600. We felt that all we had to do was replicate the CPU's to get improved performance. This later proved to be an incorrect solution to the problem. A single Instruction Processing Unit with multiple ALU's proved to be the correct solution.

As the year progressed, we realized that we had to increase our software staff and Gary Boswell was the first person hired (Appendix C). His first task was to look into the Fortran on the TIAC 870 to see if it could be modified to be usable on the ASC. As time progressed, the software effort became a major task as we sought to serve the seismic market as well as the more general purpose outside market.

Sam Smith was placed in charge of the program on 4 September 1968, replacing Harvey Cragon. At that time, the schedule was to turn power on the PPU by the end of 1968 and to turn power on the full system by the end of 1969. We were to have a Fortran compiler operational on the ASC by the end of 1970.

Buddy Dean, Wayne Winkelman, Frank Little and Gary Boswell filed a patent on the ASC operating system. This patent was issued and is one of the few pure software patents granted during this period of time.

Sometime during 1968, the development of the head-per-track disk system began. The date is not clear at this time. The development of the disk had been assigned to DSG Houston with the head development in Dallas.

The checkout of the Peripheral Processing unit started in August of 1969 on the fourth floor of the South Building (Appendix K). Drafts of the System Specification and System Control Document were completed. The total disk project was moved from Houston to Dallas and Norm Chandler was placed in charge of the effort (Appendix D).

The ASC was the first large computer system that any of us had developed, and we undertook a massive simulation effort to verify the design. This simulation required a large 360 computer dedicated to the task for long periods of time. Such a system did not exist in Dallas, but one was found at Texas A&M in College Station. Thus, an air shuttle was set up between College Station and Dallas. A team would fly down in the evening with their tapes, work during the night and return to Dallas (Appendix E). Marvin Talbott had joined the project in September 1968, and participated in the programming of the various simulators (Appendix F).

The move to Austin began in the winter of 1969/1970. We established shop on Lamar Avenue near the UT campus and the computer was moved in February 1970. This was a very nice building for a project such as the ASC, but as the tensions with students increased due to the war in Vietnam, we finally had to move out to the buildings on Highway 183 in 1972.

In February 1970, a major review of the ASC project was conducted. Presentations were given by Ray McCord, Sam Smith, Don Richmond, Milo

Backus, Grant Dove, Harvey Cragon and Ed Dawson. Ed projected a cumulative, negative cash flow of around $70M in 1974. I believe that this turned out to be very close to the actual value.

The PPU had been running reasonably well in Dallas before the move. This subsystem was made operational and the balance of the computer checkout started in mid 1970. Early in 1970, we began to have major reservations on the viability of the thin film memories we were using. It appeared that the performance and reliability would not achieve our requirements. The cost of these memories would continue to be high and the requirement for memory was growing. The memory size had been scoped based on the very small memory, highly tuned hands-on environment of SSD operations. The use of an operating system and compiler was not initially desired by SSD. As these capabilities were added, the memory needed to grow.

There was strong evidence that semiconductor memory would replace all other forms of memory then in use. Thus, a project was set up under George Quimby to investigate memory alternatives (Appendix G). Two memories were finally developed, a high speed bipolar main memory (first installed in May 1972), and a lower speed MOS attached memory (first installed in mid 1972). We were one of the earliest users of the 1K Intel 1103 device. We had a contract with a firm in Canada who had government money to get into the semiconductor business and had an agreement with Intel to produce the 1103. I remember that this firm never made it and folded in the mid 70's.

Significant problems had developed between TI and CUC on the software. There were many reasons for these problems; however, the net result was that

TI assumed control of the software development in Austin. See Appendix C for a complete account. We had started working on marketing of the ASC in 1970 and an applications group was established in Dallas to perform applications oriented programming. We received a small study contract from the Army Balistics Missile Defense Agency to study the applicability of the ASC to balistic missile defense. John Blakemore and Marvin Applewhite made contributions to this study (Appendix H).

By the fall of 1970, a 7 platter disk drive with a full complement of heads was operating in Dallas. It is interesting to note that the bearing used to support the horizontal shaft which held the platters was the same bearing used in the landing gear of a Boeing 747 aircraft. A development unit of the disk was delivered to Austin in the spring of 1971. This was followed in the second quarter with a unit having all heads and finally four units were shipped in the summer of 1971.

During the checkout of ASC 1, major reliability problems put a severe strain on both the hardware and software checkout (Appendix I and J). In conjunction with S/C, a new family of ECL circuits had been developed for the ASC. These circuits dissipated a large amount of power, and the package gave many of these problems in reliability. During the period April to August, 153 out of 378 failures were ECL circuit failures. Problems associated with the printed circuit boards accounted for 99 of the failures. Thus, 252 failures or 66 percent of the total were related to these two component groups. These were also under our control to fix. However, as 1971 was a lean year, the resources available to invest in fixing the problem were limited and the schedule continued to suffer.

An agreement was made with Royal Dutch Shell in 1971 to install a TI owned ASC in Amstelveen Holland. This agreement was the first customer agreement to use an ASC. Even though this would be a TI machine, we felt that we were starting to achieve a market recognition for the computer. In December of 1972 our efforts in marketing to ABMDA paid off, and a sale of ASC #3 was made to Charlie Vick in Huntsville. This sale was made possible by the execution of a demonstration program on the ASC at Lamar Street. This demonstration called the "Zero Order Demonstration" was a major accomplishment. We developed a real time operating system for the ASC with a BMD application running under it. The General Research Corporation of Santa Barbara, California had a radar threat environment generator running, which the ASC would interface with via telephone lines. In retrospect, it is a wonder that we ever made this system work with all the problems, but work it did, and we received the first real outside sale (Appendix H). ASC #3 was delivered to Huntsville and passed its acceptance test in August 1973.

A TI Board of Directors meeting was held in Austin in the Fall of 1971. A major demonstration of the ASC was planned, and we had just barely gotten the CPU to function as a unit, and the first disk had just been operated on-line. Doss Dunlap had been after us to show that the ASC could add two plus two and get four. This we were unable to do, but we could show the results of various diagnostic programs. Due to the fear of not being able to even run the ASC because of a component failure, several runs of the diagnostics were made and the results printed so that we would have something to show.

The move to the Highway 183 site took place in 1972. Construction on System #2, which was to go to Holland, had started in 1971 and was being assembled at the 183 site. This machine was completed and shipped by air from Dallas to Holland in November of 1972. A DC-8 and a 707 were required to transport the computer.

As System 1 was moved, it was upgraded with new circuits and boards to correct the reliability problems and became system 1A. This system was finally shut down on November 15, 1982 at 8:15 A.M. after almost 10 years of faithful service (Appendix K). System 2 was returned to Austin in 1976 and decommissioned. Many of its components were used to further upgrade 1A and other components were donated to The Computer Museum in Boston, Massachusetts.

Marketing continued during 1972 in attempts to sell the ASC to other customers. The AEC had requested bids on three machines, one each for LLL, Los Alamos and Brookhaven. We submitted a bid of $40M, CDC bid $22M, and we had lost the business. The AEC and CDC were unable to come to terms and the AEC recompeted the contract. Gary Boswell was in charge of preparing this bid, which we also lost. His account of this effort can be found in Appendix C. LLL awarded a contract to CDC for a STAR in November. The price was $6.6M.

The first outside sale of an ASC was to Geophysical Fluid Dynamics Laboratory of NOAA. This machine was ASC #4, even though #3 was sold later to ABMDA. The order was received on 13 January 1972. Appendix K describes the milestone requirements for the system integration. This machine was a four pipe configuration with a large peripheral content. Specific benchmarks were to be performed to demonstrate its performance. System integration was to be completed in June 1973. (I need to get the complete story of the GFDL machine).

In early 1973, it was clear that the marketing of the ASC was not going well. A team was assembled to conduct an in-depth study of the market for and the marketability of the ASC. In April, the report of this group was published. This report, ASC Push, covered the history, status and outlook of the project. Based on the forecasted sales, the project continued according to plan. Section I and II of this report are found in Appendix K.

In Holland, we started to run simple jobs on ASC #2 in early 1973. In the second quarter, it became obvious that the ASC was not performing as had been forecasted. The ASC was processing seismic data at a rate of .01 miles per hour while the 870 was processing at 5 miles per hour. On 17 May, a recovery team headed by Joe Watson was formed to fix the Holland problem. Reporting to Joe were Carroll Hall, Wayne Winkleman and Jim Farmer. On 22 May, Sam Smith was placed in overall charge of the ASC program, including the Services Group software effort. Ray McCord made temporary changes in the Equipment Group management, so that he could devote much of his time to the Holland problem. Joe Zimmerman became acting manager of the Equipment Group. On 1 June, Ray asked Sam to set goals for the recovery program, suggesting a first goal of 3-5 MPH and a second goal of 15-30 MPH. Sam responded with a single goal of 5 MPH by 31 August. This goal was met, and the eventual performance of the ASC was in the 100 MPH range.

The above mentioned performance improvement of 10,000 was brought about by improvements in the TIPEX operating system and improvements in the Fortran compiler. This history of the ASC is primarily from a hardware point of view and does not fully comprehend the effort which went into the software system. There were several changes in direction or goals for the software system which complicated the program. In all cases, the software staff was able to regroup and produce a working system. I believe that the Fortran compiler effort was of significant importance. A vectorizing compiler which generated very efficient code was designed, built and became operational in the mid 70's. Cray and others are just now (mid 80's) delivering compilers of equal quality. This compiler, named NX, produced optimized code but compiled at a very slow rate. This was fine for production codes but unacceptable for testing codes. Thus, a fast compiler called FX was developed in one year and completed in 1974. The goal of compiling 30,000 lines per minute was achieved.

A COBOL compiler was also developed in the 1975 period of the NRL ASC. The software story should be told in a separate volume. Appendix C, written by Gary Boswell, and Appendix F, written by Marvin Talbott, contain the best available accounts of the general purpose software development. Appendix H, by Marvin Applewhite, describes the real time software effort for ABMDA and the development of a Process Design Language.

As the dust cleared from the Holland problem, Sam asked the senior staff for critiques of the ASC program on 4 August. Sam's condensation and interpretation of the critique inputs can be found in Appendix L.

Return now to the GFDL ASC. This machine was delivered in April 1974 and passed its acceptance test at GFDL in November 1974. GFDL required a large secondary storage system for the ASC. We investigated several approaches to this requirement and selected an array of video tape recorders which were modified to store digital data. After much difficulty, this system finally passed acceptance tests in 1975 (month?). However, the system would not stand up to round-the-clock operation, and we finally negotiated a reasonable settlement with GFDL. This system represents the only complete hardware failure of the ASC project; all other problems were solved in some way. For example, the answer variability problem at GFDL was solved by redesign of the memory error checking boards and the thin film memory problem was solved by the use of semiconductor memories.

The GFDL ASC was turned off and replaced by an IBM??? in ??. (I need the story of the last days at GFDL).

The marketing of an ASC to the Naval Research Laboratory (NRL) was taking place during the Holland problem time. However, marketing was essentially stopped, and Gary Boswell and Buddy Dean went to NRL to tell them that we would not sell them an ASC even if they wanted one. NRL had some unusual requirements such as the computer supplier had to build the building to house the computer. Gary continued to work with NRL as the seismic performance improved, and we finally submitted a firm bid to NRL and won the contract. As of this date, May 1984, the NRL ASC is still performing well.

Joe Watson left TI in the summer of 1974 and Gary Boswell was made program manager of the ASC project.

I need the story of the installation and operation of the ASC at Hunstville. When was it turned off? What is the story of the Army Corps of Engineers at Vicksburg, Mississippi? Where is it now? Was it scrapped? etc.

ASC #5 was the second seismic machine and was installed in the Austin Seismic center in December 1974. A third seismic machine, ASC #6, was completed for Austin in January 1975.

The ASC Push report of 1973 had identified the need for a lower cost positioning-arm-disc for the system. This effort began in the South Building basement under the direction of Bill Norvell. Consequently, the development of a system with standard IBM plug-compatible 3330 type disc was started for ASC use. The transfer rate of these discs was being increased by a factor of nine to meet the requirements of the ASC. The basic 3330 disc had a transfer rate of 6.4 million bits per second and this was being increased to 40 million bits per second. (I do not remember how this was being done. Were we reading 8 discs in parallel, adding more electronics to each disc?) This development was finally completed in the winter of 1976 (date?) when the checkout of the first PAD took place on ASC 1A.

The last bid of an ASC was to NOAH in Colorado. This bid was lost to Cray, I believe, in the Spring of 1976 and was the end of the ASC program. The development efforts were shut off as was any new production. Most of the personnel on the program were rapidly reassigned to other programs and only a sustaining effort for the seismic machines and the NRL machine were maintained. In January the last ties with CUC were broken and all software was under TI control. Upgrades to the seismic machines continued to be made, and two of these machines are still in operation. Of the six ASC's produced, three are still in service.