the Availability Digest

www.availabilitydigest.com

Try Doing This Today

August 2013

In today's technology with gigabytes of memory packed into small laptop computers having gigahertz of speed, we forget (or never knew) what it was like in the 1960s and 1970s to develop major systems. Let's go back to those days when memory was measured in kilobytes and processor speed was measured in megahertz and see what miracles (in today's perception) were achieved. Think about what it would take to do this today.

A Payroll Services Company

MiniData Payroll Services was established in 1969. Its vision was to provide payroll services to small companies using low-cost computing. When it went into operation, its primary competitor was the soon-to-be giant ADP, which was ignoring the small-company market.

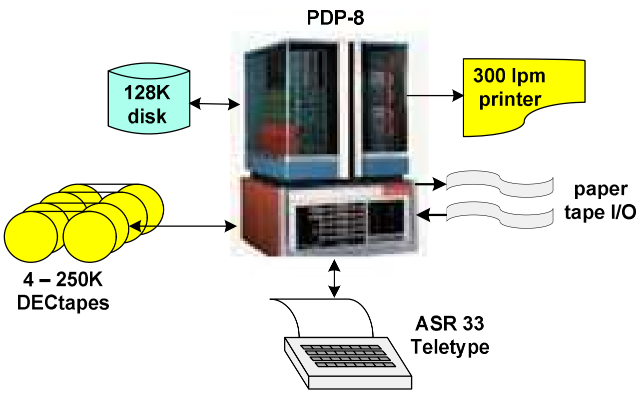

MiniData's billboard ads read:

> You pay us 9
> We'll pay 15

The message was that MiniData would do a fifteen man payroll for $9.

To accomplish this, MiniData purchased a Digital Equipment Corporation PDP-8 computer. The PDP-8 was at that time the best-selling of DEC's computers (later to be overshadowed by the PDP-11, which was the forerunner of the VAX and VMS). A basic PDP-8 cost $18,000 and came with 4K of memory with a 12-bit word length (back then, a character was six bits – ASCII had yet to make it into the computer world). Yes – four kilowords of memory, not four megawords of memory. In today's terminology, it was basically an 8K machine. Its processing speed was a blinding one megahertz – 1/1000 of today's speeds.

The I/O devices that came with the basic PDP-8 included a Model 33 ASR Teletype machine that printed at ten characters per second. It used paper tape for its primary I/O, reading paper tape at a rate of 200 characters per second and punching paper tape at 60 characters per second. The Teletype and paper tape units used 8-bit ASCII code.

To the basic PDP-8 model, MiniData added another 4K of memory (giving it 16K 6-bit bytes), 128K of disk storage, and four 250K DECtapes. The DECtapes had the interesting capability of being able to read and rewrite sectors on tape – they were effectively sequential disk storage devices. A high-speed 300 line-per-minute printer rounded out the configuration.

The PDP-8 came with virtually no operating system except for the device drivers. The MiniData payroll applications were written in SAIBOL-8 (SAI Business Oriented Language) developed by its parent, Sombers Associates, Inc. SAIBOL-8 was a COBOL-like language implemented as an interpretive compiler. Basically, each verb was a subroutine call that executed the functions of the verb.

The software architecture for the payroll applications used the lower 4K bank of memory to hold the device drivers and the SAIBOL-8 subroutines. It also included a page-swapping system for swapping 4K blocks of SAIBOL-8 routines from the disk into the upper 4K of memory.

The payroll customer database was held on the DECtapes. Customer records could be read from the DECtapes and updated by rewriting the data onto the tapes.

One problem that MiniData faced was that it could not run one payroll at a time. To do so meant breaking down the line printer for each customer's processing pass to insert the company's checks and then inserting blank paper to print its reports. To solve this problem, MiniData printed checks on blank stock. It could then print the checks for fifty customers at a time followed by the customer reports, thus essentially eliminating printer breakdown time. It did its own MICR encoding of the checks (the machine-readable numbers at the bottom of the checks) for each customer to finish the payroll.

MiniData meant with a great deal of success. By the 1980s, it was servicing hundreds of customers and graduated to dual PDP-11s. It is still in business servicing thousands of customers.

Apollo 11 – Landing a Man on the Moon

We talk a lot about continuous availability – systems that virtually never fail – as if this is a brand new concept made possible by today's technologies of fault-tolerant systems, clusters, virtualization, and active/active systems. But forty years ago, continuously available systems put men on the moon and got them back several times without a loss. The entire lunar mission, from launch to translunar coast to landing and then rendezvous and return to earth, were controlled by computers whose failure would mean the loss of the astronauts. Against all odds, all systems worked impeccably for the lunar landings.[1]

How did the technologies of the '60s achieve this near-miracle? Three main computer complexes were involved in each lunar flight – the ground computers, the onboard guidance computers, and the rocket computer.

Ground Computers

To put technology into perspective, at that time the computer industry was just in the process of converting from vacuum tubes to transistors. The five computers received by NASA in 1964 were among the first transistorized IBM 360s delivered. The 360s had a speed of 1.7 MIPS (million instructions per second) and contained one megabyte of memory. Memory was later expanded to two and then to four megabytes per computer. Five 360s were used to provide high redundancy.

The Apollo 11 IBM Computer Room

The new MVS operating system was a batch-oriented system but had to do a real-time job. Specialized communication plug boards were developed to connect to the I/O bus to receive ground tracking, onboard telemetry, and trajectory data in real time. The

---

[1] Apollo 11 – Continuous Availability, 1960's Style, *Availability Digest*; September 2009. http://www.availabilitydigest.com/public_articles/0409/apollo_11.pdf

plug boards interfaced to the computer bus via a five-way switch so that they could be connected to any of the five computers.

At the time, IBM described the six megabyte programs it developed to monitor the spacecraft's environmental data and the astronauts' biomedical data as the most complex software ever written.

Redundancy was provided via the use of three of the five IBM computers working in parallel. The other two computers were available for test and development but could be pressed into production if one of the three production computers went down.

Of the three production computers, one was the operations computer, one was a dynamic standby, and one was a static standby. The operations computer and the dynamic standby processed all data in parallel. If one went down, the other immediately and automatically took over control and became the operations computer. At this point, the static standby computer was brought into service as the dynamic standby. It was brought up-to-date by loading the last data checkpoint and then by replaying all subsequent events at high speed.

During this recovery time, operations continued unimpeded. A total system failure would have required that following a failure of one of the active computers, the other would fail before the new standby was synchronized. Even then, only a small amount of data would have been lost; and the system would quickly resynchronize with newly received data. Continuous availability was truly achieved.

Guidance Computers

The onboard Apollo Guidance Computers (AGC) controlled the Columbia Command Module and the Eagle Lunar Module. They collected flight information, displayed this to the astronauts, and communicated it to the ground. The guidance systems provided for manual control should an AGC exhibit problems.

The AGC was the first computer to use integrated circuits (ICs). All 4,100 ICs were identical three-input nor (negative or) gates. They were configured to provide flip-flops for data storage and to provide the logic of the computer. The ICs were socket-mounted and were interconnected by wire wrap.

The memory of each of the AGC lock-stepped computers comprised 2K words of random-access RAM memory and 36K of read-only ROM memory, each with a cycle time of about 12 microseconds. The word length was 16 bits. The processor ran at a speed of one megahertz and could multitask eight jobs at a time via time slots.

Modern-day PCs may be more powerful than the AGC, but the AGC did much more than plug into a printer and a router. It interfaced with ground telemetry links, radar rendezvous systems, landing altimeters, gyro compasses, optical star trackers, and propulsion systems.

The Rocket Computer

The final computer in this trilogy is the one that guided the Saturn 5 rocket at liftoff. The Launch Vehicle Digital Computer (LVDC) was embedded within a three-foot high, 21-foot diameter ring that mounted onto the Saturn V rocket. It was the nerve center for the launch vehicle.

The LVDC had a two megahertz clock rate and contained 32K, 28-bit words. For reliability, it used triple-redundant logic with voting. Even with this, calculations showed only a 99.6% reliability over 250 hours of operation. On the other hand, imagine the environment in which this computer had to operate. After less than six hours of productive activity radioing onboard measurements back to Earth and computing guidance control and engine control, the LVDC completed its life work. The third-stage Saturn V rocket engine nosed down into the ocean, and the computer ring was turned loose to orbit the sun.

Race Track Wagering at Aqueduct, Belmont, and Saratoga

The New York Racing Association (NYRA) operates some of the best known race tracks in the world – Aqueduct, Belmont, and Saratoga. Belmont is the home of the Belmont Stakes, one of the three racing events that together are known as the Triple Crown.

All race tracks, including NYRA, operate *totalizator* systems. A totalizator is the system which accumulates wagers into pools. Based on the monies accumulated in each pool, it calculates and displays the current odds that each horse in the race may come in first (win), second (place), or show (third) and displays the payoffs for the winning horses at the conclusion of the race.

Up until the early 1960s, totalizators were electromechanical relay systems. These systems were inflexible in that adding additional wagering types such as a Daily Double (choosing the first and second place horses) or a Trifecta (choosing the top three winners) was very difficult. In addition, they were not as reliable as would be desired and required continual maintenance by highly skilled technicians.

In the early 1960s, NYRA decided to build a computerized totalizator system to solve these problems. This had never been done before, so NYRA was the pioneer.[2]

This first computerized totalizator system was to perform all of the functions of its electromechanical kin. It needed to service 600 ticket issuing machines (TIMs) generating 1,200 transactions per second. Win, Place, Show, and Daily Double pools were to be supported; and the system had to be expandable to handle other types of wagering.

The availability of the totalizator system was a critical issue. There had been cases of mechanical totalizators failing just before a race, and patrons were not able to place bets on a horse that in fact won and paid big money. This had resulted in riots and significant damage to race track facilities. Therefore, totalizator failure had to be avoided at all costs. 100% uptime was paramount. The Recovery Time Objective (RTO) was zero.

In addition, the system had to be protected from data loss should it fail. The pools still had to be accounted for, and monies still had to be paid out. The Recovery Point Objective, the measure of the amount of data that could be lost due to a system failure, was also zero.

The system chosen for this task was the Honeywell H200 (yes, Honeywell was a major computer supplier back then). The H200 was a six-bit character-oriented machine, making it an unlikely system for such requirements. However, it was blazingly fast and had an advanced interrupt system for its time. It was deemed to be one of the few systems capable of handling the transaction workload.

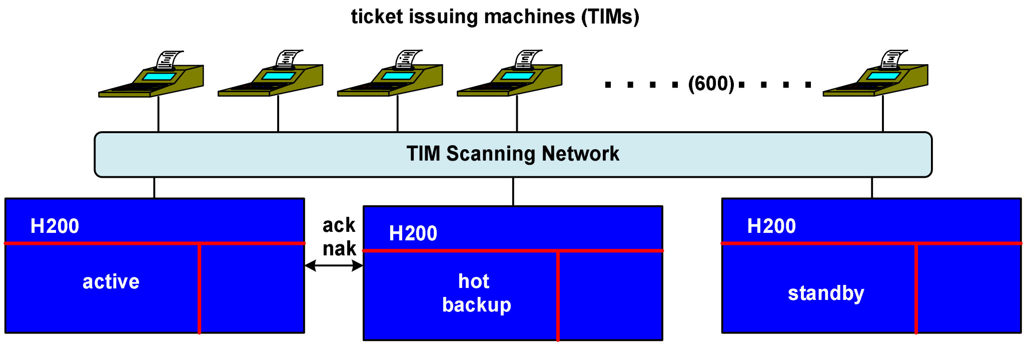

The First Computerized Totalizator

---

[2] CPA at Aqueduct, Belmont, and Saratoga Racetracks, *Availability Digest*; January 2007. http://www.availabilitydigest.com/private/0201/nyra.pdf

The system was designed to be a triplexed system (triple modular redundancy). Three identically configured H200 systems cooperated in the operation. One acted as the active system, one as a hot backup, and one as a cold standby.

The active system and the backup system were maintained in synchronism via lock-stepping at the transaction level. Should the active system fail, an operator command switched roles; and the backup system immediately took over the active role. The only delay in system recovery was the operator's decision time to switch over. Should either the active system or the backup system fail, the cold standby was put into service as the new hot backup system. It was synchronized with the current active system by transferring the current state of the pools via a memory-to-memory transfer, a process which took only a few seconds and was done concurrently with continued wagering.

To achieve the required 1,200 transaction per second processing rate, all processing of a ticket sale was done at the interrupt level.

This early, continuously available system remained in operation for almost two decades. It morphed into several later versions of very successful totalizator systems installed at hundreds of tracks around the world. Following the H200 system was one based on Digital PDP-8s, then Tandem systems, followed by Digital PDP-11s and VAXs. In later years, both the Tandem systems and the VAX systems morphed into state lottery systems.

The Chicago Board of Trade

The Chicago Board of Trade (CBOT) is the world's oldest futures and options exchange. A wide variety of commodities are traded via open outcry and electronic trading.

In the 1960s, all trading was via open outcry. Traders stood in pits, an octagonal structure of steps leading to a flat floor. They indicated their intent to buy or sell at specified prices by hand signals and yelling. When two traders agreed to make a trade, they pointed at each other. Each wrote the trade on a slip of paper and threw it into the center of the pit.

A Chicago Board of Trade Open Outcry Pit

Clerks gathered the trade slips from the floor of the pit and put them on conveyer belts. The first stations that the conveyer belts passed were Morse-code operators who tapped out the trade to men walking around a massive raised chalk board. These men listened to the Morse-encoded trades and wrote the trades on the chalk board for the trading room to see. From the Morse code operator stations, the slips then moved to a back room for after-hours settlement.

This system came to an end after railroads stopped using Morse code to signal between stations. This was when the source of Morse-code-trained operators began to dry up.

The CBOT decided to install Ferranti display boards to replace the old chalk boards. These were mechanical displays that used rotating numerical drums to display trade prices. "Pulpits" were placed at the edge of each pit, and operators observing the pit activity entered the trades into terminals in the pulpits. These messages were sent to a PDP-8 that controlled the Ferranti display boards. The terminals printed a hard copy of all trade activity. Slips were still used. They were gathered and sent to the back office for trade settlement.

Three PDP-8s were used. One was engaged in active trade reporting, and the others were standbys. Each was an 8K machine (16K 6-bit characters).

Much of the open outcry trading has since been replaced with electronic trading. However, open outcry trading in the pits is still in use by the CBOT.

Summary

In our current high technology, it is difficult to see how complex, real-time systems requiring continuous availability could be built with only a few thousand bytes of memory and with clock speeds measured in the megahertz range. These case studies reflect the brilliance of early computer professionals, who built major systems with such minimum resources – a brilliance that has led to today's plethora of computing capabilities.