DOWNLOAD PDF ILLIAC IV, THE FIRST SUPERCOMPUTER

Chapter 1 : | Timeline of Computer History | Computer History Museum

The ILLIAC IV was the first massively parallel computer. The system was originally designed to have bit floating point units (FPUs) and four central processing units (CPUs) able to process 1 billion operations per second. Due to budget constraints, only a single "quadrant" with 64 FPUs and a single CPU was built.

However, the technical report was a major influence on computing in the s, and was used as a blueprint for many other computers , including two at the University of Illinois, which were both completed before Princeton finished Johnniac. Taub , Donald B. Gillies , Daniel Slotnick played a key role in the computer design s. It was the first von Neumann architecture computer built and owned by an American university. It was put into service on September 22, By it had gained more computing power than all computers in Bell Labs combined. Data was represented in bit words , of which 1, could be stored in the main memory, and 12, on drum memory. At its inception in it was times faster than competing machines of that day. It became operational in , two years later than expected. The core memory access time was 1. A "fast buffer" was also provided for storage of short loops and intermediate results similar in concept to what is now called cache. The "fast buffer" access time was 0. Floating-point numbers used a format with 7 bits of exponent power of 4 and 45 bits of mantissa. Instructions were either 26 bits or 13 bits long, allowing packing of up to 4 instructions per memory word. The pipelined functional units were called advanced control, delayed control, and interplay. The computer used Muller speed-independent circuitry i. Muller C-Element for a portion of the control circuitry. In Donald B. Later it was used on biological images. The machine was destroyed in a fire, caused by a Variac shorting on one of the wooden-top benches, in Key to the design as conceived by Daniel Slotnick, the director of the project, was fairly high parallelism with up to processors, used to allow the machine to work on large data sets in what would later be known as array processing. The machine was to have 4 quadrants. Several years into the project, TI backed out and said that they could not produce the LSI chips at the contracted price. This required a complete redesign using medium scale integrated circuits,leading to large delays and greatly increasing costs. This also led to scaling the system back from four quadrants to a single quadrant, owing to the fact that the MSI version was going to be many times larger than the LSI version would have been. For the PEs what should have been chips about 1 inch in diameter were now roughly 6 by 10 inches. Space, power and air conditioning not to mention budget did not allow for a four quadrant machine. There could be instructions being sent from the CU on the wires to the PEs at any time. The power supplies for the machine were so large that it required designing a single tongue fork lift to remove and reinstall the power supply. The power supply buss bars on the machine spanned distances greater than three feet, and were octopus-like in design. Thick copper, the busses were coated in epoxy that often cracked resulting in shorts and an array of other issues. Illiac IV was designed in fact to be a "back end processor" to a B The cost overruns caused by not getting the LSI chips and other design screw ups by Burroughs the control unit was built with positive logic and the PEs with negative logic, etc. First, that the project had been secretly created on campus. When this claim proved to be false, the focus shifted to the role of Universities in secret military research. Slotnick was not in favor of running classified programs on the machine. ARPA wanted the machine room encased in copper to prevent off site snooping of classified data. Slotnick refused to do that. He went further and insisted that all research performed on Illiac IV would be published. If the machine had been installed in Urbana this would have been the case. One was that Slotnick was concerned that the physical presence of the machine on campus might attract violence on the part of student radicals. The machine was never delivered to Illinois, arriving in In , when the first and only quadrant was operational at NASA, it was 13 times faster than any other machine operating at the time. The development team was led by Professor David Kuck. This SMP symmetric multiprocessing system embodied advances in interconnection networks, control unit support of parallelism, optimizing compilers and parallel algorithms and applications. It was a node Linux cluster, with each node having two processors. Trusted ILLIAC nodes contained onboard FPGAs to enable smart compilers and programming models, system assessment and validation, configurable trust mechanisms, automated fault management, on-line adaptation, and numerous other configurable trust

frameworks. The nodes each had access to 8 GB memory on a 6.

Chapter 2 : history of comp timeline | Timetoast timelines

The Illiac IV was the first large scale array computer. As the foreÂ- runner of today's advanced computers, it brought whole classes of scientific computations into the realm of practicality.

According to the Jan. In return for the computer, government agents would be able to use it for military research. Illiac IV was to be the most powerful computer on the face of the earth at that time. The protests reached a boiling point on May 9, , in a day of "Illiaction. Eckert and John Mauchly. He wins another in for co-developing the theory of superconductivity. He would become professor of physics and electrical engineering at the University in The Digital Computer Lab is organized. It was ten feet long, two feet wide, eight and one-half feet high, contained 2, vacuum tubes, and weighed five tons. ILLIAC, the first computer built and owned entirely by an educational institution, becomes operational. It was used by Lajaren Hiller, director of the Experimental Music Studio, to compose and play the Illiac Suite, the first computer-composed composition. UI faculty publish what is believed to be the first journal article in behavioral and social sciences involving a computer. Studies are underway of advances such as transistors, parallel operation, high-speed circuitry, and improved logic to better the usefulness, speed, and reliability of computers. Robertson, an electrical engineer who was an expert in error-checking systems, pioneers basic techniques of efficient binary division. The SRT division algorithm, now found both in hardware and software implementations of the divide instruction and widely used in the most powerful microprocessors, is named after D. Sweeney, Robertson, and T. The campus got an IBM , which was used in the design of research instruments like high-energy particle accelerators and radio telescopes. ACM Computing Reviews says of the machine, " ILLIAC II, at its conception in the mids, represents, together with some other independent design projects of the same period, the spearhead and breakthrough into a new generation of machines. The machine was to analyze bubble chamber photographs of high energy particle events. Due to a building fire, it was never finished. Gillies discovered three Mersenne prime numbers in the course of checking out ILLIAC II, including the largest then known prime number, , which is over 3, digits, putting him in the Guiness Book of Records for a time. Slotnick had worked under John von Neumann at Princeton. It was also the first to employ ECL Emitter-Coupled Logic integrated circuits and multilayer up to twelve layers circuit boards on a large scale. Most notable was its use of semiconductor memory. Undergraduate degree program in computer science is established in the College of Engineering.

Chapter 3 : The Illiac IV: The First Supercomputer - R.M. Hord - Google Books

ILLIAC IV was a SIMD computer (single instruction, multiple data) and it marked the first use of circuit card design automation outside IBM. It was also the first to employ ECL (Emitter-Coupled Logic) integrated circuits and multilayer (up to twelve layers) circuit boards on a large scale.

If you are a fan of the novel and movie Why pick such an odd location? Urbana is hardly a household name unless you know the Chicago area well. But Urbana has a place in real-life computer history. Sometimes you learn more from failure than you do successes and at least one of the ILLIAC series is the poster child for that. The Urbana story starts in the early s. In it, Von Neumann proposed changes to EDVAC that would make it a stored program computer â€" that is, a computer that treats data and instructions the same. The exact number is hard to pin down since many machines reused older machines, but however you count, there were at least a dozen. The Army took delivery of the new machine in It used over 2, vacuum tubes mostly 6J6s and s and had 1K of bit words using Williams tubes. As part of the contract, the Army agreed to fund the construction of an identical machine for the University to keep. Programs written on one could be run on the other. That seems like no big deal today, but in there were these machines were so large and expensive that only one of each was built. This made it one of the first, if not the first, computer to be operated remotely. The computers weighed in between 3, and 5, pounds and used hexadecimal. This made some kind of sense based on the teleprinters used, apparently. Other systems, used other letters, apparently depending on what was handy for their hardware. Usually, CPUs have a master clock that each circuit obeys. So one unit would finish its work and signal the next unit to start. This early version of PLATO serviced a single user, although version two could accommodate two simultaneous users. Errors Were Common Thinking of HAL as a computer mind that went wrong is a proposition very different today than it would have been at the time. Our modern hardware has unbelievably low error rates. But errors and hardware failures were particularly common in the s and had to be worked around by the computer operators. Periodic tests under stress conditions would identify flaws, hopefully before they affected normal operations. A later project sped the readers up to reduce that number to under 8 minutes. Keep in mind, this was about 40K bits of data, so imagine what it would take to fill up a modern flash drive. Just as we tinker with our machines, these computers often grew, and one reference reports that by when ILLIAC retired, it had 5, words of memory along with storage on a magnetic drum, pictured below. The design for this machine started in and it came online in This was a transistorized machine with 8K words of core memory. The word size was 52 bits. This machine used asynchronous modules and pipelining to get a form of parallel execution. Thanks to these enhancements, the machine was about times faster than earlier machines of the day. The machine would have 4 quadrants and each quadrant would have a CPU a control unit and processing elements PEs. When complete, the 4 CPUs would have the ability to each operate on 64 pieces of data at one time. But the project was plagued by problems from the start. The new design had to use off-the-shelf ICs leading to much larger cards with more cooling and power requirements. This led to more additional costs. Another cost problem occurred when they realized the control unit used positive logic and the processing elements used negative logic. Despite its problems â€" and frequent downtime due to power supply problems, the machine was still 13 times faster than other machines in In , it was connected to the ARPANet â€" the predecessor of the Internet â€" and became the first network-accessible supercomputer. The use of large-scale ECL integrated circuits, thin-film memory, and specialized disk drives that held a whopping 80 megabytes on a inch disk, were all risky items for their day. It even had plans for a write-once laser memory system that could store a terabit of data on a drum carrying a piece of polyester. Keep in mind that â€" at the time â€" 20 logic gates in a package was considered large-scale. This slowed the clock speed from 25 MHz to 16 MHz, pushed the project over budget and ran the schedule out 2 years. TI successfully produced the chips a year later and sold them commercially, but by that time the team was too far down the redesign path. The board size increase took up the room slated for the thin film memory. So you make the cabinets bigger, right? That caused problems with signal propagation and distribution. So the team changed to newfangled semiconductor memories. An Ambitious Project Cut Down to

Size By , running massively over budget and schedule, they decided to only do one quadrant. That cut the potential speed of the machine significantly. Operationally, the machine was a nightmare to bring up. PCBs were cracked, wirewrap terminals would short, card contacts were prone to oxidation, and the card socket plastic softened over time. The ICs were very sensitive to humidity, and there were many other issues. By the computer could run programs, but the results were not always correct. In , Ames started a 4-month effort to fix things. They replaced over , resistors, rewired long propagation delay lines, cleaned up the power supplies, and reduced the clock speed from 16 MHz to 13 MHz, only half the speed of the original 25 MHz design. Now the machine was able to do its work reliably at least some of the time. Ames also replaced the original Burroughs control computer with a PDP The operating technique reflected the need to avoid hardware errors. If they passed, they would run your job. If the diagnostics passed, you got the results of your job and the bill for your time. The performance was lackluster at first, although a custom version of Fortran helped programmers take better advantage of the hardware. On problems that were amenable to parallel execution, the machine was the fastest in the world until According to the same source, an IBM would need nearly 3 seconds for the same task. Semiconductor memory was the way of the future. So was large-scale IC integration. The larger boards required more grounding and pushed the team to layer boards which would be easy now, but were difficult to produce in those days. Even mass storage using write-once laser media was something that would become important years later. Or you can dig into the gory details. It is largely considered a failure but it was useful and it did lead the way in a number of important technologies. It also overpromised and under delivered.

Chapter 4 : CiteSeerX â€" Citation Query R.: The Illiac IV the first supercomputer

The ILLIAC II was the first transistorized and pipelined supercomputer built by the University of Illinois. ILLIAC II and The IBM Stretch were two competing projects to build 1st-generation transistorized supercomputers.

A formal design did not start until , when Slotnick was working at Westinghouse Electric and arranged development funding under a US Air Force contract. With funding from Advanced Research Projects Agency ARPA , they began the design of a newer concept with bit processors instead of a single machine with 1, 1-bit processors. While the machine was being built at Burroughs , the university began building a new facility to house it. It is also credited with being the first large computer to use solid-state memory, as well as the most complex computer built to date, with over 1 million gates. This drum had 80 tracks so two words could be read at a time, and each track stored 1, bits. If the bits of a word were written serially to a single track, instead of in parallel across 40 tracks, then the data could be fed into a bit-serial computer directly from the drum bit-by-bit. The drum would still have multiple tracks and heads, but instead of gathering up a word and sending it to a single ALU, in this concept the data on each track would be read a bit at a time and sent into parallel ALUs. This would be a word-parallel, bit-serial computer. By this time, for scientific computing at least, tubes and drums had been replaced by transistors and core memory. The idea of parallel processors working on different streams of data from a drum no longer had the same obvious appeal. Nevertheless, further consideration showed that parallel machines could still offer significant performance in some applications; Slotnick and a colleague, John Cocke, wrote a paper on the concept in During this period, some consideration was given to more complex PE designs, becoming a bit parallel system that would be organized in a by arrangement. A single PE using this design was built in As the design work continued, the primary sponsor within the US Department of Defense was killed in an accident and no further funding was forthcoming. Westinghouse management considered it too risky, and shut down the team. Slotnik left Westinghouse attempting to find venture capital to continue the project, but failed. Illinois had been designing and building large computers for the U. Development started in , and a first-pass design was completed in One of these, RGR, was used for communicating data to neighbouring PEs, moving one "hop" per clock cycle. This allowed the system to work on different problems when the data was too small to demand the entire PE array. Generally, these languages provided support for loading arrays of data "across" the PEs to be executed in parallel, and some even supported the unwinding of loops into array operations. Seventeen responses were received in July, seven responded, and of these three were selected. In August ,[b] eight-month contracts were offered to RCA , Burroughs and Univac to bid on the construction of the machine. Both offered new technical advances that made their bid the most interesting. Burroughs was offering to build a new and much faster version of thin-film memory which would improve performance. They would also provide a Burroughs B mainframe to act as a front-end controller, loading data from secondary storage and performing other housekeeping tasks. The more complex internal wiring was causing crosstalk in the circuitry, and they asked for another year to fix the problems. TI was able to get the pin design working after just over another year, and began offering them on the market before ILLIAC was complete. Attempts to increase the size of the cabinets to make room for the memory caused serious problems with signal propagation. On 6 January , The Daily Illini , the student newspaper, claimed that the computer would be used to design nuclear weapons. However, he also grew increasingly concerned that the machine would be subject to attack by the more radical student groups. The machine was finally delivered to Ames in April , and installed in the Central Computer Facility in building N This caused further delays in bringing the machine online. When the machine first arrived, it could not be made to work. It suffered from all sorts of problems from cracking PCBs, to bad resistors , to the packaging of the TI ICs being highly sensitive to humidity. These issues were slowly addressed, and by the summer of the first programs were able to be run on the system although the results were highly questionable. At the end of this process, the system was finally working properly. On problems that could be parallelized the machine was still the fastest in the world, outperforming the CDC by two to six times, and it is generally credited as the fastest machine in the world until One control unit and one processing element chassis from

the machine is now on display at the Computer History Museum in Mountain View, less than a mile from its operational site. It was widely considered a failure even by those who worked on it; one stated simply that "any impartial observer has to regard Illiac IV as a failure in a technical sense. As Slotnik himself later put it: Delighted that the overall objectives came out well in the end. Slotnick received a lot of criticism when he chose Fairchild Semiconductor to produce the memory ICs, as at the time the production line was an empty room and the design existed only on paper. As Slotnick would later comment, "Fairchild did a magnificent job of pulling our chestnuts out of the fire. The Fairchild memories were superb and their reliability to this day is just incredibly good. As their complexity grew, the PCBs had to add more and more layers in order to avoid growing larger. Eventually, they reached layers deep, which proved to be well beyond the capabilities of draftsmen. The design was ultimately completed using new automated design tools provided by a subcontractor, and the complete design required two years of computer time on a Burroughs mainframe. This was a major step forward in computer aided design , and by the mids such tools were commonplace. It was the better understanding of parallelism on ILLIAC that led to the improved compilers and programs that could take advantage of these designs. Similar to the ILLIAC in some ways, these processor designs loaded up many data elements into a single custom processor instead of a large number of specialized ones. There was more than a little "backlash" against the ILLIAC design as a result, and for some time the supercomputer market looked on massively parallel designs with disdain, even when they were successful. As Seymour Cray famously quipped, "If you were plowing a field, which would you rather use? Two strong oxen or chickens? Each PU was identical to all of the others, so they could be replaced or reordered as required. The system had a short instruction pipeline and implemented instruction look ahead. The fourth, R, was used to broadcast or receive data from the other PEs. Both the PEs and CU could use load and store operations to access the disk system. For this reason, the CU could not be used to coordinate actions, instead, the entire system was clock-synchronous with all operations in the PEs guaranteed to take the same amount of time no matter what the operands were. That way the CU could be sure that the operations were complete without having to wait for results or status codes. The eight-away connections allowed faster transport when the data needed to travel between more distant PEs. This allowed, for example, the same instruction stream to work on data that was not aligned in the same locations in different PEs. The common example would be an array of data that was loaded into different locations in the PEMs, which could then be made uniform by setting the index in the different PEs. Normally, when the CPU completes processing an instruction, the program counter PC is incremented by one word and the next instruction is read. This process is interrupted by branches , which causes the PC to jump to one of two locations depending on a test, like whether a given memory address holds a non-zero value. Since those values are private to the PE, the following instructions would need to be loaded based on a value only the PE knew. Logical tests did not change the PC, instead, they set "mode bits" that told the PE whether or not to run the next arithmetic instruction. To use this system, the program would be written so that one of the two possible instruction streams followed the test, and ended with an instruction to invert the bits. Code for the second branch would then follow, ending with an instruction to set all the bits to 1. When it reached the end of that code, the mode operator instruction would flip the mode bits, and from then on that PE would ignore further instructions. This would continue until it reached the end of the code for the second branch, where the mode reset instruction would turn the PE back on. For instance, in the case when the PE was operating in bit mode, one "side" of the PE might have the test come out true while the other side was false.

Chapter 5 : ILLIAC was HAL â€™s Granddaddy

Note: Citations are based on reference standards. However, formatting rules can vary widely between applications and fields of interest or study. The specific requirements or preferences of your reviewing publisher, classroom teacher, institution or organization should be applied.

This drum had 80 tracks so two words could be read at a time, and each track stored 1, bits. If the bits of a word were written serially to a single track, instead of in parallel across 40 tracks, then the data could be fed into a bit-serial computer directly from the drum bit-by-bit. The drum would still have multiple tracks and heads, but instead of gathering up a word and sending it to a single ALU, in this concept the data on each track would be read a bit at a time and sent into parallel ALUs. This would be a word-parallel, bit-serial computer. By this time, for scientific computing at least, tubes and drums had been replaced by transistors and core memory. The idea of parallel processors working on different streams of data from a drum no longer had the same obvious appeal. Nevertheless, further consideration showed that parallel machines could still offer significant performance in some applications; Slotnick and a colleague, John Cocke, wrote a paper on the concept in During this period, some consideration was given to more complex PE designs, becoming a bit parallel system that would be organized in a by arrangement. A single PE using this design was built in As the design work continued, the primary sponsor within the US Department of Defense was killed in an accident and no further funding was forthcoming. Westinghouse management considered it too risky, and shut down the team. Slotnik left Westinghouse attempting to find venture capital to continue the project, but failed. Illinois had been designing and building large computers for the U. Development started in , and a first-pass design was completed in One of these, RGR, was used for communicating data to neighbouring PEs, moving one "hop" per clock cycle. This allowed the system to work on different problems when the data was too small to demand the entire PE array. Generally, these languages provided support for loading arrays of data "across" the PEs to be executed in parallel, and some even supported the unwinding of loops into array operations. Seventeen responses were received in July, seven responded, and of these three were selected. In August , [b] eight-month contracts were offered to RCA , Burroughs and Univac to bid on the construction of the machine. Both offered new technical advances that made their bid the most interesting. Burroughs was offering to build a new and much faster version of thin-film memory which would improve performance. They would also provide a Burroughs B mainframe to act as a front-end controller, loading data from secondary storage and performing other housekeeping tasks. The more complex internal wiring was causing crosstalk in the circuitry, and they asked for another year to fix the problems. TI was able to get the pin design working after just over another year, and began offering them on the market before ILLIAC was complete. Attempts to increase the size of the cabinets to make room for the memory caused serious problems with signal propagation. On 6 January , The Daily Illini , the student newspaper, claimed that the computer would be used to design nuclear weapons. However, he also grew increasingly concerned that the machine would be subject to attack by the more radical student groups. The machine was finally delivered to Ames in April , and installed in the Central Computer Facility in building N This caused further delays in bringing the machine online. When the machine first arrived, it could not be made to work. It suffered from all sorts of problems from cracking PCBs, to bad resistors , to the packaging of the TI ICs being highly sensitive to humidity. These issues were slowly addressed, and by the summer of the first programs were able to be run on the system although the results were highly questionable. At the end of this process, the system was finally working properly. On problems that could be parallelized the machine was still the fastest in the world, outperforming the CDC by two to six times, and it is generally credited as the fastest machine in the world until One control unit and one processing element chassis from the machine is now on display at the Computer History Museum in Mountain View, less than a mile from its operational site. It was widely considered a failure even by those who worked on it; one stated simply that "any impartial observer has to regard Illiac IV as a failure in a technical sense. As Slotnik himself later put it: Delighted that the overall objectives came out well in the end. Slotnick received a lot of criticism when he chose Fairchild Semiconductor to produce the memory ICs, as at the time the

production line was an empty room and the design existed only on paper. As Slotnick would later comment, "Fairchild did a magnificent job of pulling our chestnuts out of the fire. The Fairchild memories were superb and their reliability to this day is just incredibly good. As their complexity grew, the PCBs had to add more and more layers in order to avoid growing larger. Eventually, they reached layers deep, which proved to be well beyond the capabilities of draftsmen. The design was ultimately completed using new automated design tools provided by a subcontractor, and the complete design required two years of computer time on a Burroughs mainframe. This was a major step forward in computer aided design , and by the mids such tools were commonplace. It was the better understanding of parallelism on ILLIAC that led to the improved compilers and programs that could take advantage of these designs. Similar to the ILLIAC in some ways, these processor designs loaded up many data elements into a single custom processor instead of a large number of specialized ones. There was more than a little "backlash" against the ILLIAC design as a result, and for some time the supercomputer market looked on massively parallel designs with disdain, even when they were successful. As Seymour Cray famously quipped, "If you were plowing a field, which would you rather use? Two strong oxen or chickens? Each PU was identical to all of the others, so they could be replaced or reordered as required. The system had a short instruction pipeline and implemented instruction look ahead. The fourth, R, was used to broadcast or receive data from the other PEs. Both the PEs and CU could use load and store operations to access the disk system. For this reason, the CU could not be used to coordinate actions, instead, the entire system was clock-synchronous with all operations in the PEs guaranteed to take the same amount of time no matter what the operands were. That way the CU could be sure that the operations were complete without having to wait for results or status codes. The eight-away connections allowed faster transport when the data needed to travel between more distant PEs. This allowed, for example, the same instruction stream to work on data that was not aligned in the same locations in different PEs. The common example would be an array of data that was loaded into different locations in the PEMs, which could then be made uniform by setting the index in the different PEs. Normally, when the CPU completes processing an instruction, the program counter PC is incremented by one word and the next instruction is read. This process is interrupted by branches , which causes the PC to jump to one of two locations depending on a test, like whether a given memory address holds a non-zero value. Since those values are private to the PE, the following instructions would need to be loaded based on a value only the PE knew. Logical tests did not change the PC, instead, they set "mode bits" that told the PE whether or not to run the next arithmetic instruction. To use this system, the program would be written so that one of the two possible instruction streams followed the test, and ended with an instruction to invert the bits. Code for the second branch would then follow, ending with an instruction to set all the bits to 1. When it reached the end of that code, the mode operator instruction would flip the mode bits, and from then on that PE would ignore further instructions. This would continue until it reached the end of the code for the second branch, where the mode reset instruction would turn the PE back on. For instance, in the case when the PE was operating in bit mode, one "side" of the PE might have the test come out true while the other side was false.

Chapter 6 : ILLIAC - Wikipedia

The Illiac IV was the first large scale array computer. As the fore runner of today's advanced computers, it brought whole classes of scientific computations into the realm of practicality.

Chapter 7 : ILLIAC was HAL â€™s Granddaddy | Hackaday

The Illiac IV was the first large scale array computer. Today the Illiac IV continues to service large-scale scientific aoolication areas includ- ing computational fluid dynamics, seismic stress wave.

Chapter 8 : iLLiac IV supercomputer - retroComputingTasmania

DOWNLOAD PDF ILLIAC IV, THE FIRST SUPERCOMPUTER

The Illiac IV, the first "Internet" connected supercomputer. The Illiac IV was host 15 on the ARPAnet (ARPAnet map showing Illiac IV), and most users were connecting from Host 16 to submit jobs to the Illiac IV.

Chapter 9 : ILLIAC IV - Wikipedia

The ILLIAC IV was one of the first attempts to build a massively parallel racedaydvl.com of a series of research machines (the ILLIACs from the University of Illinois), the ILLIAC IV design featured fairly high parallelism with up to processors, used to allow the machine to work on large data sets in what would later be known as vector processing.