Chapter +9: Apollo-11 Lunar Lander

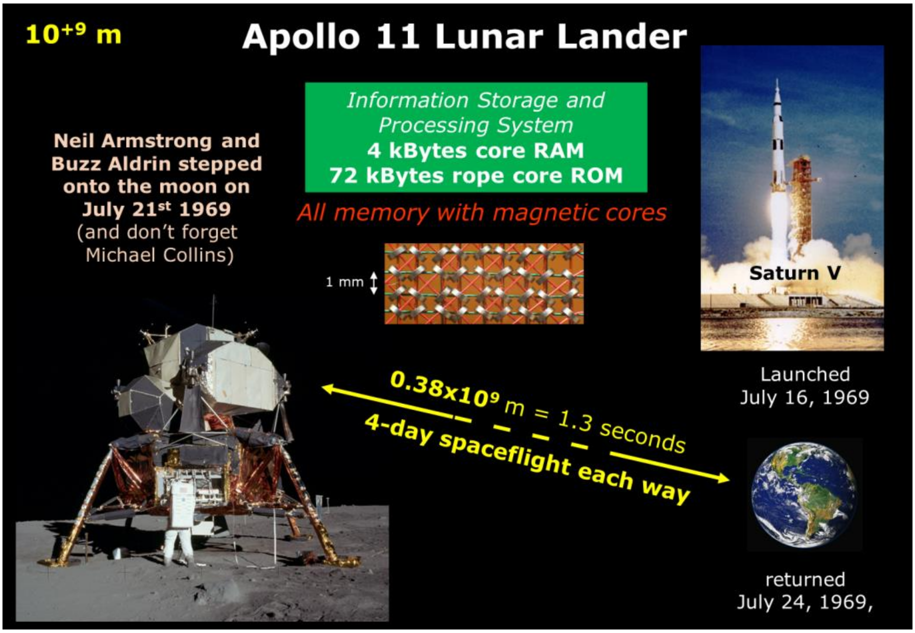

10^{+9} m

Apollo 11 Lunar Lander

Information Storage and Processing System
4 kBytes core RAM
72 kBytes rope core ROM

All memory with magnetic cores

Neil Armstrong and Buzz Aldrin stepped onto the moon on July 21st 1969
(and don't forget Michael Collins)

0.38×10^9 m = 1.3 seconds
4-day spaceflight each way

Launched
July 16, 1969

returned
July 24, 1969,

Merritt Island is the largest of the barrier islands on Florida's Atlantic coast. It is a huge wetland area designated as a National Wildlife Refuge and is home to numerous species of reptiles and mammals and birds. For these resident creatures, July 16, 1969, probably began much as any other day. Even the larger animals - alligators and manatees and great blue herons - probably showed little interest in the huge black and white cylindrical object the busy humans had been putting together. It was certainly big – towering 110 meters (380 feet) into the sky and about 10 meters (33 feet) across. But at 9:32 AM that morning, the strange object could no longer be ignored. The tranquil setting was suddenly interrupted by an ear-splitting roar that tore the air. White-hot flames shot from the bottom of the cylinder and a massive pall of smoke rose into the air. Five huge engines had been ignited and were burning through some 13 tonnes[1] of fuel every single second. The engines generated about 3,500 tonnes-force (7.6 million pounds-force) of thrust – more than enough to hurl the 3,000 tonne

---

[1] The tonne (= 1000 kg) is a convenient unit for large rockets. It is very similar in size to the English ton (2240 pounds) and about 10% bigger than the US ton (2000 pounds). A mass of one tonne weighs one tonne-force (9807 Newtons) on the Earth's surface

(6.5 million pound) cylinder up into the sky on its plume of fire and for it to vanish out into the beyond.

Four nights later, the crescent Moon was lighting up the western sky. Some of the animals in the wildlife refuge probably glanced at it. A lot of humans in Florida would have been staring at it – almost in disbelief. Some might even have brought out a telescope to peer more closely. The Sea of Tranquility could be seen as a dark blotch close to the terminator[2] nearly half way down. A higher-power telescope would show it as a relatively flat area but pock-marked with craters and surrounded by mountains. What no Earth-based telescope could possibly show was, towards the southern end of the Sea of Tranquility, an ugly angular metallic object about 10 meters in size (30 feet) that was resting on the surface. Nor at about 11 pm (EST) on that Sunday night, could anyone on Earth possibly have seen two tiny figures emerge - two intrepid explorers actually standing and walking on the surface of that waxing crescent Moon - two intrepid explorers themselves staring up in awe at a huge blue and white jewel set into a jet black sky.

The Space Race

The story behind the Apollo-11 mission and that Saturn-V launch from the Kennedy Space Center started in earnest about 12 years earlier. In October 1957, the author, then 6-years old, recalls listening to some strange beeps re-broadcast on BBC radio. They seemed to generate some excitement in the household. The beeps were coming from Sputnik, the first artificial satellite of Earth. It was launched in a high-inclination (65^o to the equator) elliptical orbit and consequently was widely visible[3] over much of the Earth. Its beeps on 20 MHz and 40 MHz could be readily picked up on short-wave by radio amateurs as it passed overhead. Sputnik was launched by the Soviet Union, the cold-war rival of the United States. It was launched using a modified Inter-Continental Ballistic Missile (ICBM).

The Soviets quickly realized the propaganda success achieved with Sputnik. Over the next few years, this first artificial satellite (Sputnik, Oct. 1957) was quickly followed up by other successes: the first

---

[2] The terminator is the line that separates the sunlit side from the dark side.

[3] What most people saw was actually the much-larger much-brighter (magnitude 1) final rocket stage that also ended up in orbit. Sputnik was magnitude 6 and barely visible.

animal in space (Laika, Nov. 1957); the first spacecraft to hit the Moon (Luna-2, Sep. 1959); the first pictures of the far side of the Moon (Luna-3, Oct. 1959); the first man in space (Yuri Gagarin, Apr. 1961); the first woman in space (Valentina Tereshkova, June 1963); the first spacewalk (Alexey Leonov, Mar. 1965); and the first soft-landing (unmanned) on the Moon (Luna-9, Feb. 1966). The overwhelming superiority of Soviet technology and thus the Soviet communist system could not be made more clear.

Sputnik and the rapid string of Soviet successes shocked the Western powers. This was at the height of the cold war and of the nuclear arms race. Both sides were vying to demonstrate the biggest atomic bomb explosions and, after that, the biggest hydrogen bomb explosions (Tsar Bomba at 50 Megatons TNT). The Soviet Union’s success in rocketry and spaceflight and *payload delivery* was viewed as an existential threat by the United States. And so started the “Space Race”.

In response to Sputnik, the US immediately brought forward its launch plans. The US Navy’s Vanguard rocket was launched in December, 1957, carrying a tiny 1.3 kg (2.9 lbs) satellite. It rose about a meter (3 feet) from the launch-pad, stalled, fell back, and exploded. The satellite landed in some bushes nearby and began transmitting signals, prompting one American journalist to remark, "Why doesn't somebody go out there, find it, and shoot it?"[4] The American press called it “Kaputnik”. The Russian delegate at the United Nations offered aid "under the Soviet program of technical assistance to backwards nations". This highly publicized and humiliating failure let to the responsibility being passed to Werner von Braun’s rocket team with the US Army.

The first successful US orbital launch occurred on January 31, 1958 using a rocket derived from the Redstone missile. Redstone was in turn a direct descendant of the German V-2 liquid-fuel rocket, developed in Germany during the war under Von Braun’s guidance. The satellite called Explorer-1 was sent into an elliptical orbit extending from 360 km to 2,500 km altitude (220 to 1600 miles). The satellite instrumentation was designed and built by Dr. James Van Allen who had included a Geiger-Müller tube to detect radiation. The tube rapidly saturated whenever the satellite reached certain altitudes confirming the existence of the eponymous radiation belts that Van Allen had theorized.

---

[4] https://en.wikipedia.org/wiki/Vanguard_(rocket)

On July 29, 1958, President Eisenhower signed into law a bill creating the National Aeronautics and Space Administration (NASA) to oversee all the civilian activities and "To provide for research into problems of flight within and outside the earth's atmosphere". Von Braun was appointed director of the newly formed Marshall Space Flight Center in Huntsville, Alabama – apparently only after he had insisted on being allowed to continue development of the Saturn rocket program.

On April 12th, 1961, Yuri Gagarin became the first human to orbit the Earth. Five days later, on April 17th, was the "Bay of Pigs" debacle, a failed invasion of Cuba sponsored by the US to overthrow Fidel Castro, the young communist leader. President John Kennedy, who previously had been lukewarm to the proposal to land people on the Moon, saw the need to save face and to re-establish US leadership. On May 25, 1961, Kennedy sent a "Special Message to Congress on Urgent National Needs":

> "... I believe that this nation should commit itself to achieving the goal, before this decade is out, of landing a man on the Moon and returning him safely to the Earth. No single space project in this period will be more impressive to mankind, or more important for the long-range exploration of space, and none will be so difficult or expensive to accomplish."

The goal was set and the race was on!

Getting There and Getting Back

There were several options originally considered for getting humans to the Moon and safely back. The more direct approaches generally required the development of a more powerful rocket and/or were not feasible within the allowed time-frame. The best strategy seemed to be to carefully stage the mission so that the minimum mass (hence minimum energy) was used at each stage. This was the approach selected. What was worrisome, however, was that this approach was very reliant on the skill of the three astronauts in handling the multiple docking and undocking procedures of the several vehicle components that made up the spacecraft. The diagram below (Apollo-11 Flight Profile) and the following list outline the mission sequence and the critical docking and rendezvous procedures:

1) Saturn-V stages 1 and 2 and a partial burn of stage 3 send all the required hardware and three crew members to low Earth orbit.

2) A second (last) burn (the trans-lunar injection burn) from the Saturn 3rd-stage sends the entire assembly en route to the Moon.
3) The Command/Service module detaches, rotates end-to-end, and then goes back to dock with and extract the Lunar Module which was contained within the 3rd-stage faring.
4) The combined Command/Service module plus Lunar module coast towards the moon for 3 days (as does the now separated 3rd-stage).
5) Drawn by the Moon's gravity, both spacecraft swing around the back of the Moon. The Service module engine is fired to bring the Command/Service/Lunar modules into Lunar orbit (the discarded 3rd–stage rounds the moon and gets thrown into a Solar orbit).
6) Once the desired circular Lunar orbit is established, two crew members enter the Lunar module which is then separated from the Command/Service module where the third crew member remains.
7) The lunar module fires its descent stage engines to break out of orbit and descend to the surface. The engines are used to slow the descent and make a carefully controlled landing on the surface.
8) After completing the activities on the surface, the engine on the Lunar module ascent stage fires separating it from the lower descent portion (which remains on the Moon). The ascent stage rises back into Lunar orbit close to the orbiting Command/Service module with the lone crew-member.
9) The Command/Service module and the Lunar Module ascent stage dock together. The crew members are reunited in the Command module. The Lunar ascent stage is detached and abandoned.
10) The Service module engine is fired again to break out from Lunar orbit and to set the Command/Service module on its 3-day journey back to Earth.
11) As the Earth is approached, the Service module is jettisoned (it burns up in the atmosphere) and the Command module is oriented with its heat-shield towards the Earth - ready for a fiery re-entry.

12) The high return velocity (11 km/s or 7 mi/s)[5] of the Command module is absorbed by its thick heat-shield in the Earth's atmosphere. Finally a large parachute opens and guides the Command module to a not-so-gentle (35 km/hr or 22 mi/hr) splash-down in the Pacific Ocean ready to be picked up by the US Navy.

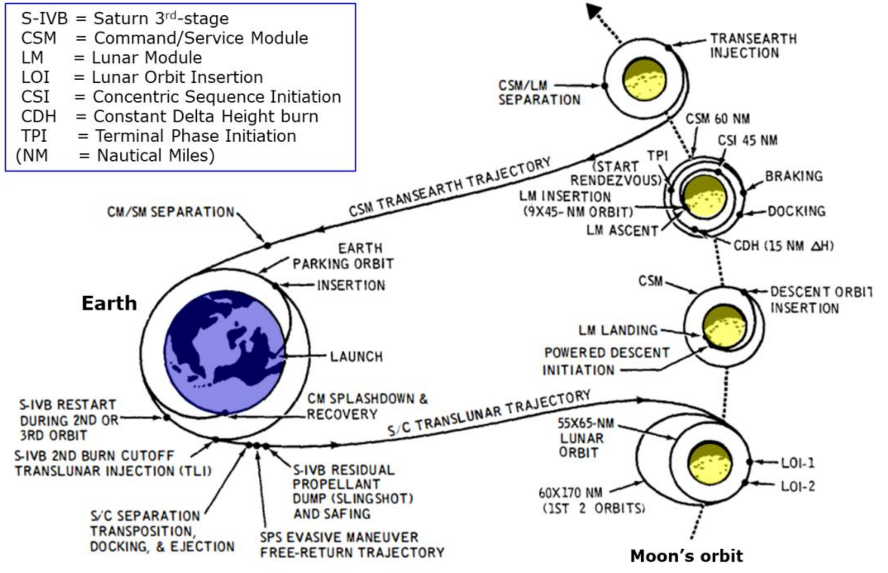

Apollo-11 Flight Profile: 8 days 3 hours 18 mins from launch to splash-down
https://www.hq.nasa.gov/alsj/a11/A11_MissionOpReport.pdf

The workhorse for the mission was the Saturn-V rocket (that huge black and white cylinder referred to in the opening paragraph). It was launched from the Kennedy Space Center on Merritt Island. The Saturn-V rocket remains the tallest and heaviest and most powerful rocket that has ever been built (as of 2018). It was capable of hefting an incredible 140 tonnes (310,000 pounds) into low Earth orbit[6]. The man behind this rocket design was Werner von Braun, a German scientist who came to work in the US at the end of World War II. The primary contractors doing the construction were Boeing, North

---

[5] The Moon is relatively far away so the return speed is close to Earth escape speed which is sqrt(2) higher than low Earth orbit velocity. Twice the kinetic energy must be dissipated compared with a return from low Earth orbit.

[6] Space-X Falcon-Heavy, which is currently (2018) the most powerful rocket, can move 64 tonnes to low Earth orbit. The three first-stage rockets can be reused – a big advantage. https://en.wikipedia.org/wiki/Falcon_Heavy

American Aviation, Douglas Aircraft Company, and IBM. The Saturn-V was a three-stage rocket. The massive first-stage weighted 2,300 tonnes of which 93%[7] was fuel (RP-1[8] and liquid oxygen). It burned for 168 seconds. The second stage weighed 480 tonnes of which 94% was fuel (liquid hydrogen and liquid oxygen). It burned for 360 seconds. The third stage weighed 119 tonnes of which 92% was liquid hydrogen and liquid oxygen. This stage was designed to make two burns for a total 500 seconds. Altogether only fifteen Saturn-V rockets were produced (nothing was re-usable). Four rockets were used in testing and then ten were used for the manned Apollo flights and the last was used to launch SkyLab. There was never a serious failure in the launch vehicle.

There were just four manned Saturn-V/Apollo test-flights prior to the Moon-landing. Apollo-7 was the first manned orbital flight and tested out the Command/Service Module (CSM). Apollo-8 put the first humans in orbit around the Moon and confirmed the ability of the Command module to re-enter Earth's atmosphere at very high-speed. Apollo-9 was kept in Earth orbit to test out all the new components including the Lunar module and the space-suit for Extra Vehicular Activity. This also provided the opportunity to practice the complex docking procedures with the Lunar module. Apollo-10 on May 18th 1969 was a complete dry-run for the Lunar landing. The Lunar module was separated and brought to within 16 km (10 miles) of the Moon's surface. The crew then had to rendezvous and transfer back to the Command Module, all while in Lunar orbit. Finally, just two months later, was the real-thing, *Apollo-11*.

The Apollo-11 Moon landing was highly publicized in real-time. Fortunately, as we know, everything went to plan and the astronauts returned to a tumultuous reception. America had regained its leadership is a most dramatic and decisive manner. The Soviet Union's plans to land their cosmonauts on the moon first were beset by numerous failures and delays. Indeed, they never really acknowledged that there ever had been a race. The Apollo program went on to complete six more missions to the Moon. Only one mission (Apollo-13) suffered a major malfunction in the Service module that caused the

---

[7] For reference, about 96% of a modern soda can is soda. The other 4% is aluminum.

[8] RP-1 is a grade of Kerosene that is carefully refined so that it can do double-duty as fuel and as the coolant for the combustion chamber surfaces. The light and heavy fractions are removed so it consists mainly of hydrocarbons around weight C_{12}.

moon landing to be aborted. Heroic efforts did enable the crew to return safely[9].

The Lunar Module

The Lunar module (sometimes known affectionately as the LEM, short for Lunar Excursion Module) was a two-stage vehicle designed to get the astronauts down to the surface of the Moon, support a short period of exploration on the surface, and then return the astronauts to Lunar orbit and the rendezvous with the Command/Service module. The ascent stage was stacked on top of the descent stage and the two stages had separate engines and separate fuel supplies. The Lunar module was obviously designed with a strict emphasis on minimizing weight. Nevertheless, the combined mass of the two stages reached a substantial 15 tonnes fully fueled.

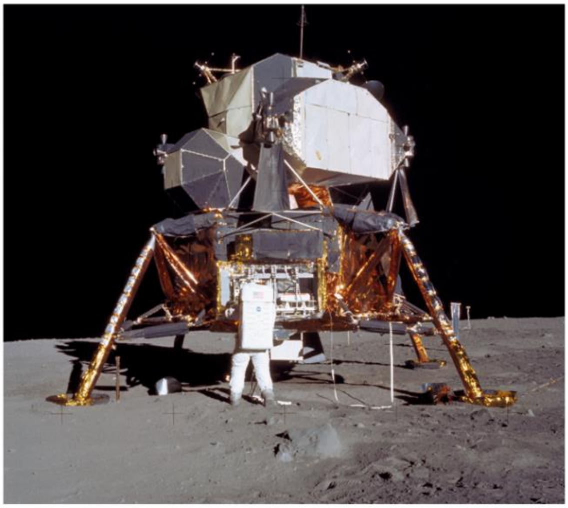

Apollo-11 Lunar Module in the Sea of Tranquility.

Neil Armstrong took this photograph of "Buzz" Aldrin deploying the experiments package on July 20th, 1969.

The upper part of the Lunar module provides life-support for the two crew members and contains the fuel and engines required for the ascent back to lunar orbit and the waiting Command/Service module.

The lower part of the Lunar module contains the fuel and engines used during the descent and landing phase. Today it is still standing on the Moon's surface.

https://moon.nasa.gov/resources/188/view-apollo-11-lunar-module-as-it-rested-on-lunar-surface/

The descent stage by itself had a mass of 10.3 tonnes of which 8.2 tonnes was fuel. The fuel was aerozine-50 (a more stable variation on hydrazine, N_2O_4) and nitrogen tetroxide (N_2O_4) oxidizer. This combination is hypergolic meaning, reassuringly, that it reliably ignites spontaneously as soon as the two components are introduced into the combustion chamber of the rocket engine. The descent engine was mounted on gimbals and could be throttled back to as low as 10% of the maximum 4.6 tonnes-force[10] - a necessity for controlled landing.

---

[9] https://www.nasa.gov/mission_pages/apollo/missions/apollo13.html

[10] This is more than enough to lift 15 tonnes mass in Lunar gravity (~17% of Earth's)

The descent stage formed an octagonal structure 4.2 meters (14 feet) across and 1.7 m (5.6 feet) tall. Four landing legs folded out and absorbed the landing forces and held the stage well clear of the surface. The descent stage also had storage space for the lunar equipment and experiments (and, on later missions, the lunar rover). Having served its purpose, the descent stage together with most of the experimental equipment and the cameras were all abandoned on the Moon when the ascent stage was launched.

The ascent stage was a very irregularly shaped unit mounted on top of the descent stage. It was much lighter with a mass of 4.7 tonnes of which 2.4 tonnes was fuel (same type as the descent stage). The ascent engine produced a fixed constant thrust of 1.6 tonnes-force (3,500 lbf or 16,000 N). The tiny crew compartment had an accessible volume of 4.5 cubic meters (160 cubic feet). It had no seats, but hammocks could be strung up if the crew wanted to rest or sleep. The cabin was pressurized with 100% oxygen at 33 kPa (4.8 psi or 30% atmospheric pressure). Essentially the astronauts were living in a very thin-skinned aluminum balloon.

The ascent stage had a small hatch on the side of the ascent module to provide (quite difficult) access to the Lunar surface and another hatch and docking mechanism at the top for the crew to return to the Command/Service module. Attitude control was provided by 16 hypergolic thrusters (440 N, 46 kg-force, 100 pounds-force each) grouped in fours around the periphery of the ascent module. The ascent stage with its guidance computer was responsible for all the control functions for either the ascent stage alone or for the combined vehicle during descent. It was acknowledged early in the program that the spacecraft would need to be flown by computer. This was especially true for the Lunar landing phase.

The Lunar module computer

The Apollo Guidance Computer performed the role of ‘flight computer’ meaning that its key responsibility was to control and navigate the spacecraft[11]. Its essential job was to develop and continuously integrate the spacecraft’s twelve-element state vector to accurately describe its attitude, position, velocity, and acceleration in space. The position and attitude (orientation) were initialized and frequently updated based on sextant fixes on stars, the sun, and the

---

[11] https://tcf.pages.tcnj.edu/files/2013/12/Apollo-Guidance-Computer-2009.pdf

planets (including pre-selected landmarks on Earth) and also from inputs from gyroscopes and accelerometers in the digital autopilot. The radio link back to Earth also allowed for distances (transmission delay) and velocities (Doppler-shift) and accelerations (estimated gravitational-fields) to be accurately calculated and relayed back to the spacecraft. With little human intervention, the computer was able to autonomously navigate from the Earth to the Moon, to control the Lunar landing, to control the ascent from the Moon, and to control the rendezvous with the Command module. In fact the complexity of the tasks made the computer an essential component of the mission.

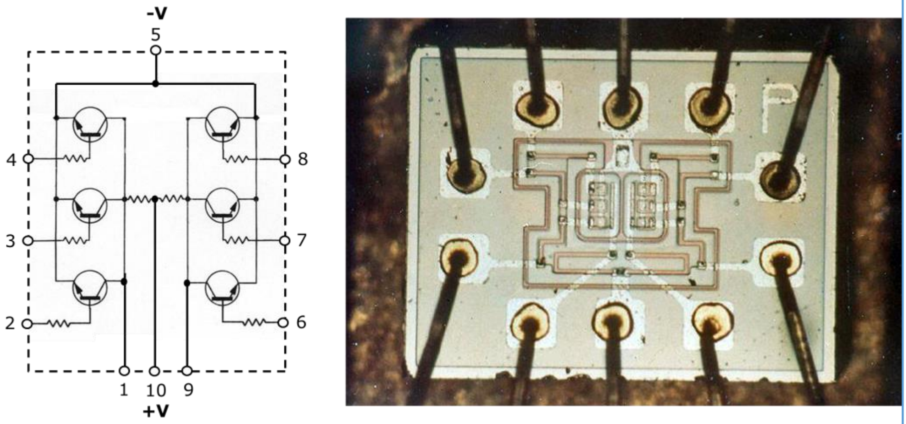

Dual 3-input NOR gate integrated circuit (IC) implemented in resistor-transistor logic. This is the only IC used in the Apollo-11 Guidance Computer and is used throughout. The IC uses six transistors - modern ICs contain over 10 billion). These are bipolar transistors which predate the FETs introduced in Chapter +12.
https://en.wikipedia.org/wiki/Apollo_Guidance_Computer#/media/File:Agc_nor2.jpg

Because of the need for proven reliability and the need for a design-freeze early in the mission, the technology behind of all the components of the Apollo mission dates back to the early 1960s. This is very true of the Apollo flight computer which was one of the very first computers to use integrated-circuits (IC). The IC had been invented just a few years earlier in 1958 or 1959 (depending on definition). The Apollo Guidance Computer[12] used 2,800 identical ICs manufactured by Fairchild. Each IC was a dual three-input NOR gate implemented in resistor-transistor logic (several resistors and just six transistors on the chip). The NOR gates (an OR function with the output inverted) were arranged to perform all the logic functions in the

---

[12] E. Hall, “General Design Characteristics of the Apollo Guidance Computer”, R-410, MIT Instrumentation Laboratory, May 1963
http://klabs.org/history/history_docs/mit_docs/1009.pdf

computer. All the arithmetic operations (including multiply and divide instructions) were performed using a single adder together with 16 ‘high-speed’ registers - all built from the same dual NOR gates. The computations were driven by a two-phase clock running at a modest 1 MHz. The ICs were interconnected with wire-wrap done by hand and then the entire assembly was potted in epoxy for mechanical stability. The design was done at the MIT Instrumentation Laboratory with Raytheon manufacturing the units. The computer had a mass of 32 kg (70 pounds) and fitted into a 61×32×17 cm (24×12.5×6.5 inch) hermetically-sealed aluminum box.

There were two of these Guidance Computers on the Apollo mission. One was in the Command Module and one in the Lunar module. The computers supported the astronauts in navigation and control of the spacecraft. The human interface consisted of a numeric keypad (like a simple calculator) for input and a digital display for output. Spacecraft commands were parsed into a verb (action) and a noun (object) with a two-digit code for each. These were entered on the keypad. The digital display was capable of showing three five-digit signed numbers (no decimal point). These would typically show vectors such as the space craft orientation or a required velocity change (delta-V)[13]. Each digit was formed using a seven segment display driven by an electromechanical stepping relay producing a distinctive clicking as the display was updated. Per a request from the astronauts, all the quantities were in imperial (US) units and were translated from the metric units that were used for calculations internally.

The word-length used in the computer was 16-bits including one parity bit. Unusual for its time, the computer had an asynchronous real-time operating system with an interrupt-driven scheduler. The software, or rather ‘firmware’ since it was very much hard-wired into the core rope read-only memory (see below), included a number of ‘interpreted’ instructions that were implemented as short subroutines. This greatly saved on the lines of code required, since coding space in the read-only memory was at a premium (36k words). The ‘interpreted’ instructions included double-precision arithmetic, trigonometric functions, vector operations, and even a matrix times vector instruction – all essential operations for navigation and control. The team that

---

[13] https://en.wikipedia.org/wiki/Apollo_Guidance_Computer

created the code for Apollo was led by Margaret Hamilton[14] at the Massachusetts Institute of Technology Instrumentation Laboratory.

Margaret Hamilton, a computer scientist who led the Software Engineering Division of the MIT Instrumentation Laboratory.

This iconic photograph was taken in 1969. The giant stack of paper beside her are the listings of the computer code that she and her team at MIT produced for the Apollo Program.

The 'firmware', responsible for all the control and navigation of the spacecraft through its various maneuvers, was painstakingly 'woven' into 72 kiloBytes of core rope memory.

http://news.mit.edu/2016/scene-at-mit-margaret-hamilton-apollo-code-0817

The basic hardware and software architecture revolved around a 3-bit instruction code and 12-bit address code inherited from the preliminary design. Various tricks were employed to expand the instruction set and expand the address space. The Apollo-11 version (called Block-II) had 34 distinct instructions available including the interpreted instructions. The memory (data storage) was expanded to 2 kilowords of re-writable random access memory (RAM), and 36 kilowords of fixed read-only memory (ROM). The RAM was implemented with conventional *core memory* and the ROM was implemented with what was known as *core rope memory*. Both memories are based on magnetic technology and magnetic cores but they could not be more different in operation.

Magnetic Core Memory

Magnetic Core Memory[15] was the predominant form of random-access memory for computers for about 20 years between about 1955 and 1975. Information is stored by virtue of whether magnetization (magnetic flux) is circulating clockwise or counter-clockwise round a

---

[14] http://news.mit.edu/2016/scene-at-mit-margaret-hamilton-apollo-code-0817
https://www.youtube.com/watch?v=4sKY6_nBLG0&feature=youtu.be&t=170

[15] https://en.wikipedia.org/wiki/Magnetic-core_memory

small ferrite core. A key enabler for the technology was the invention in 1951 of coincident-current addressing by Jay Forrester[16] a very prolific American computer engineer and inventor at Massachusetts Institute of Technology (MIT). Forrester also created computer architectures for dealing with the destructive readout that was an unfortunate inherent feature of the technology[17].

Magnetic Core Memory

Illustrating a single magnetic core from an array of similar cores. Each core stores a single bit of data by virtue of the direction in which the core is magnetized.

In coincident-current addressing, half the required current is sent down the x-line and half down the y-line. Only the core at the intersection has enough current to switch.

Changes in magnetic state are picked up as voltage pulses on the 'sense' line.

The 'inhibit' line works with the x-lines and y-lines during the write process.

(In later versions the 'sense' and 'inhibit' functions used just one wire)

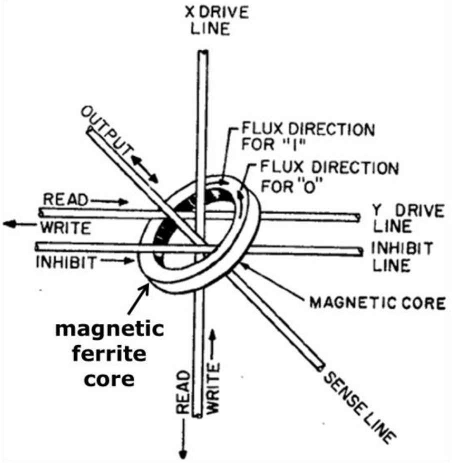

DIGITAL COMPUTER BASICS, Naval Education and Training, Command Rate Training Manual, NAVEDTRA 10088-B, Revised 1978 (*from a copy owned by Computer History Museum*)

Magnetic core memory was formed from an array of ferrite cores (tiny magnetic toroids or doughnuts) threaded by copper wires. The 'easy' (low-energy) states are for the magnetic flux to flow in a circle around the core. Passing enough current through the wire in one direction or the other direction would cause the magnetization to settle in the clockwise or counter-clockwise direction thus writing a '0' or '1'. Magnetization is difficult to sense electrically, but changes in magnetization are easy to sense since the changes produce small voltage pulses or spikes (Faraday's law). The readout strategy was to write or set the core to '1' and then check (on a separate 'sense' wire) whether or not a voltage pulse had occurred. If there was no pulse then the initial state of the core must have already been set (written) to '1'. If there was a pulse, the magnetization direction must have flipped, so the initial state of the core must have been written to '0'. Of course

---

[16] https://www.nytimes.com/2016/11/18/technology/jay-forrester-dead.html

[17] Destructive Read: reading the information was accomplished by setting the bit into a known state and detecting whether or not there was a change of state. After reading, all the bits were therefore set and the original information was destroyed.

the end result after reading was that all the cores ended up set to ‘one’ and the original information was destroyed.

The idea behind the coincident-current architecture is that there exists a relatively well-defined minimum current below which the core magnetization cannot be reversed or flipped. Above that threshold current, the magnetization reliably flips. So, if two wires are threaded through the core, it can be arranged that if current flows in just one of the two wires (either), nothing will happen. But if current flows in both of the wires simultaneously, the total current will be sufficient to set (or reset) the core reliably. This feature greatly simplifies addressing. If a rectangular array of cores (at a 45 degree angle[18]) is threaded by a first series of wires along one axis and a second series of wires along the other axis, a single individual core can be addressed by energizing a single wire along one axis and a single wire along the other axis. The only core that will respond is the one that is threaded by both energized wires where they cross each other in the array.

Core Memory from the Apollo-11 guidance computer

1024 ferrite cores are arranged in a 32 x 32 array to provide one kilobit of storage. Sixteen such planes are stacked to provide one ‘kiloword’.

There are five wires threading each core. The usual number is three or four depending on whether or not the inhibit and sense lines are multplexed. Two appear to be address lines and three appear to thread all the cores in the plane suggesting they are sense or inhibit lines.

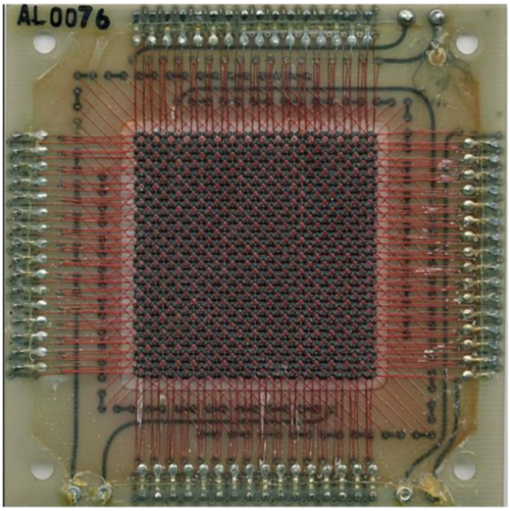

Courtesy Johnathan H. Ward
https://commons.wikimedia.org/wiki/File:Apollo_1024_bit_core_memory_module.jpg

For example, for 12 address lines, as in the Apollo computer, this provides an address space of 4096 words. Normally this would involve demultiplexing from the 12 address lines to select 1 of 4096 individual addresses. With the coincident-current approach, we can demultiplex

---

[18] The 45 degree angle makes it easier to thread the wires from both the address axes. An alternating angle of plus or minus 45 degrees makes each core orthogonal to its four immediate neighbors thus reducing interference.

6 address lines to select 1 of 64 along one axis of the array of cores and demultiplex the other 6 address lines to select 1 of 64 along the other axis – a much simpler task. The net result is the selection of 1 core in a 64 × 64 array - that is 1 of 4096.

Coincident current techniques greatly simplified hardware at a time when transistors and vacuum tubes (valves) were vastly more expensive than threading twice as many wires. If fact each core was threaded by two further lines: a sense-line and an inhibit-line. The sense-line, as the name implies, was connected to a sensitive amplifier to detect the pulse indicating that the magnetization in the selected core had switched (or not switched). The use of a separate line avoided the problem of trying to look for small pulses in the presence of the relatively large voltages present on the two select lines. The inhibit line, as its name implies, prevented the writing of the core that was being addressed. A current in the inhibit line opposed the two coincident currents selecting that core. The sense and inhibit-lines could follow exactly the same route through the cores, but, again in the early days of core memory it was easier to run two wires rather than try to make one wire perform two functions. Later designs used three-wires per core with the functions of sense and inhibit being multiplexed onto a single line.

Core memory was widely used throughout the computer industry for many years. The name still lingers in phrases like “core-dump”. Typically cores were around a 1-millimeter diameter (0.04 inches) with an inside diameter of about ½-millimeter (0.02 inches). Core-size shrank over the years. Threading even three wires successfully through the tiny hole in each core was a huge challenge. The fabrication process was initially highly manual and was only ever partially automated. Key advantages of core memory for the Apollo program were that it was non-volatile and that it had no steady-state power consumption. In the 1970’s, core memory was gradually displaced by semiconductor memory (with battery backup if a nonvolatile memory was required).

Core Rope Memory

Core Rope Memory does use magnetic cores, but in a completely different manner. *No information is stored magnetically in the cores*. The nonlinear properties of the magnetic cores are used in creating the addressing mechanism and also in forming the memory output. The data in the memory is actually woven permanently into the wiring during the fabrication of the ‘rope’. This is achieved by virtue of whether a particular wire *does or does not thread*[19] a particular core.

In core rope memory, there is one core for each address (not for each bit). That one core is selected by being the only core that is not held saturated by one or more currents flowing in the address lines (or their complements) that thread the core. The operation of selecting the only core that has all zero currents as inputs is equivalent to a logical NOR function. By threading each core with either an address line or its complement and by going through all possible binary combinations of address lines or complements, it is possible to select and assign each address uniquely to just one core that is not held in a saturated state. The single core that is not held saturated is free to respond to currents in a set/reset line, all the other cores will not respond. The magnetization in the selected core can be toggled back and forth by currents in the set/reset line, the other cores cannot.

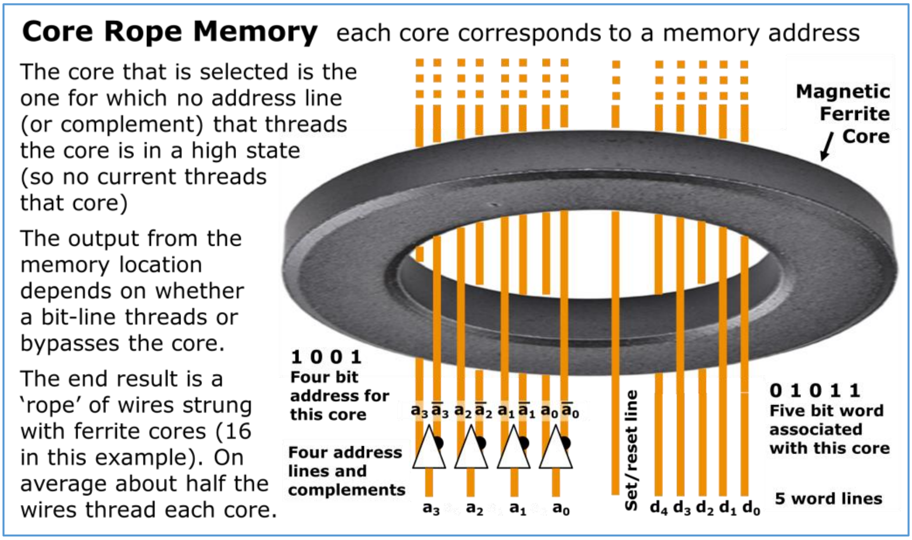

Core Rope Memory each core corresponds to a memory address

The core that is selected is the one for which no address line (or complement) that threads the core is in a high state (so no current threads that core)

The output from the memory location depends on whether a bit-line threads or bypasses the core.

The end result is a ‘rope’ of wires strung with ferrite cores (16 in this example). On average about half the wires thread each core.

As the magnetization in the core is toggled back and forth, it will create a small voltage pulse in any additional lines that thread the core. There can be many of these additional lines, each of which will show a small pulse when the core is toggled. That is how the data is stored.

---

[19] “Thread” means the wire (and current) passes through the hole in the core and creates a magnetomotive force that drives magnetic flux around the circular core. Conversely if the flux circulating in the core changes suddenly, a small voltage spike is seen on all the wires threading the core. If the wire does not thread the core and just bypasses the core, nothing happens in either writing or reading.

For the 5-bit word, shown in the example, three of the lines thread the core and will show a pulse, two lines will not. The data read back from that location will be ‘01011’. The data is manually written into the ‘rope’ by deciding which data bit lines will thread a particular core (address) and which will not.

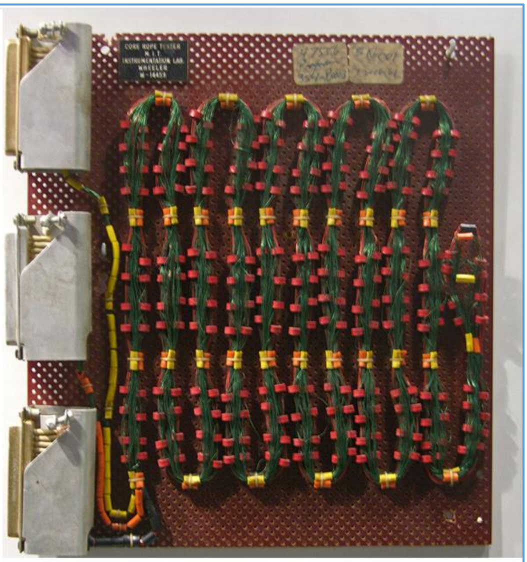

Core Rope Memory
(test sample from Apollo program)

Rope memory with 256 cores (addresses) strung along rope.

Eight address lines and their complements (16 green lines) thread (or bypass) each core to define the binary address and uniquely select the core.

Similarly, the data lines (orange) thread or bypass each core to define the data-bits in each word (the number of bits is unclear from this picture).

The rope and its memory contents are created manually. The rope is folded in half then packed in serpentine fashion onto the board.

https://commons.wikimedia.org/wiki/File:Apollo_guidiance_computer_ferrit_core_memory.jpg

Note in the picture above that the board has no transistors (or vacuum tubes) and there are no power lines supplying it. It is inherently a low-power device – the cores only consume power when reading. The address demultiplexing and the merging onto the output data-lines are both done with ferrite cores and the memory instructions are literally physically woven[20] into the device. The result is a very high reliability device that is immune to radiation (the Apollo missions had to fly through the Van Allen radiation belts) and electrical overload (Apollo-12 was struck by lightning on launch). Beyond physical damage to the board, the memory instructions were safe and the computer could always restart itself.

All this is history now. Neil Armstrong and Buzz Aldrin successfully landed the Lunar module on the Sea of Tranquility on July 20, 1969. Twenty-four hours later they had rejoined Michael Collins in the

---

[20] https://spectrum.ieee.org/tech-history/space-age/software-as-hardware-apollos-rope-memory

Command module orbiting the Moon. Three days later they were being hoisted up out of the Pacific Ocean into a Sea King helicopter and to a rapturous welcome on the deck of the USS Hornet.

The Lunar landing itself was not without excitement. During the descent phase, about a mile above the surface, the guidance computer suddenly issued several unexpected "1202" and "1201" program alarms. Mission Control in Houston, Texas, quickly assessed that it was safe to “Go-ahead” rather than hit the dreaded “Abort” button. The alarms were indicating that the computer was overloaded with interrupts. These were coming from the rendezvous radar (that was not actually needed at that time) and arose from a bug that was dependent on how the hardware randomly powered up. The software design correctly prioritized tasks as it was intended to do and the mission was never in jeopardy and was able to continue safely.

The bravery of those three individuals, Armstrong, Aldrin, & Collins, cannot be overstated – venturing into the void in a tin-can on top of a massive bomb. So much was unknown at that time and there were so many things that could have gone wrong, any one of which could have been catastrophic. All of the Lunar manned missions were conducted without loss of life, though with some close calls (Apollo-13). This amazing feat was a huge tribute to the care and skill of the entire Apollo team. However, we must never forget that three astronauts, Grissom, White, and Chaffee, did tragically lose their lives in 1967 when a fire broke out in an early version of the Command module being tested at Cape Kennedy[21].

A mere twelve years elapsed between those first beeps from the orbiting Sputnik and those live television broadcasts showing the first humans walking around on the Moon’s surface. And that was all 50 years ago. Not a single person has been back there since those heady days. The Moon is a most barren and desolate world with nothing at all to attract future visitors[22].

That perhaps has been the biggest disappointment in exploring the Solar system. We have certainly discovered a myriad fascinating ways that a world can be put together, but none that is even remotely habitable. There is no “new world” out there beckoning to be explored

---

[21] https://en.wikipedia.org/wiki/Apollo_1

[22] https://www.youtube.com/watch?v=DMdhQsHbWTs

and populated. Our human civilization blossomed into existence just a few thousand years ago. It flourishes today, but it is not clear that even as a species we will last more than another few hundred years before destroying our own habitat. This bright blue marble is so very precious.

Further Reading:

Alan Ritsko (dir.), *To the Moon*, a NOVA Production for WGBH Boston, 1999 (DVD)

APOLLO 11 (AS-506) Mission Operation Report No. M-932-69-11, June 24, 1969 https://www.hq.nasa.gov/alsj/a11/A11_MissionOpReport.pdf

David Mindell, *Digital Apollo: Human and Machine in Spaceflight*, MIT Press, Cambridge, Massachusetts, 2008

Brian Troutwine, “The Charming Genius of the Apollo Guidance Computer”, Domain Driven Design Europe, Jan. 2016, Brussels, https://www.youtube.com/watch?v=d1nz7vgyUh8

Eldon C. Hall, “General Design Characteristics of the Apollo Guidance Computer”, Report R-410, MIT Instrumentation Laboratory, May 1963, http://klabs.org/history/history_docs/mit_docs/1009.pdf

Thomas Kelley, *Moon Lander: How we developed the Apollo Lunar Module*, Smithsonian Institution Press, Washington and London, 2001

From the Earth to the Moon, HBO Miniseries, 2005 (DVD)

(see the tutorial box below in blue)

The Tyranny of the Rocket Equation[23]

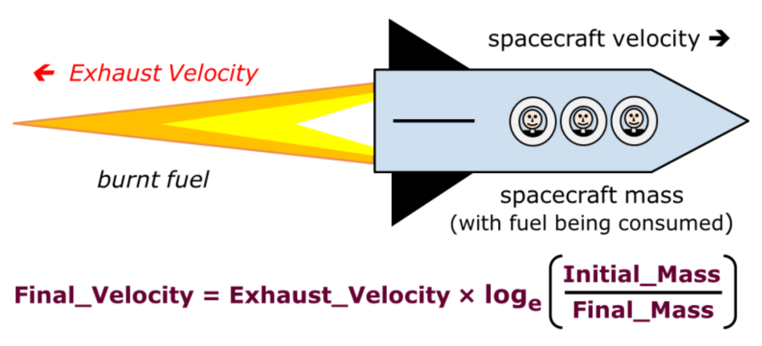

The only practical way of getting into space is to use rockets. These work by shedding mass at high velocity (exhaust) in one direction in order to accelerate the rocket in the opposite direction. Unfortunately, over most of the journey, you are accelerating not only the spacecraft but also a lot of fuel plus the structure that contains it. Conservation of momentum gives us the Tsiolkovsky rocket equation[24], written above.

Spacecraft must reach very high velocities to be useful: low Earth orbit requires 7.8 km/s (17,500 mph) while escape velocity is 11.2 km/s (25,050 mph). These numbers are much higher than the velocities obtained from any chemical reactions. Solid propellants can reach about 2.5 km/s (5,600 mph) and a rifle bullet gets to about half of that. Liquid hydrogen plus oxygen is about as good as it gets for fuel – the water-vapor exhaust travels about 4.4 km/s (9,800 mph). All these numbers are smaller than orbital velocity. The giant canon in Jules Verne's 1865 novel is strictly impossible - with or without passengers.

The $\log_e(M_{starting}/M_{final})$ factor provides the necessary amplification. A liquid-fuel rocket that is 90% propellant and 10% structure (fuel tanks payload, etc.) can reach a final velocity of 4.4 km/s × $\log_e(10)$ = 10.1 km/s. This is more than enough to get into orbit though not enough to escape Earth's gravity. It has been noted recently that aliens living on a super-earth (say twice Earth's diameter with twice the escape speed) would probably never escape their home planet. The ratio $M_{initial}/M_{final}$ would need to reach about 100. It would seem impossible to build a chemical rocket light enough yet able to carry sufficient (~99%) fuel. (*note:* making a rocket with several stages softens the tyranny a bit)

---

[23] Don Pettit, "The Tyranny of the Rocket Equation", NASA, May 2012 https://www.nasa.gov/mission_pages/station/expeditions/expedition30/tryanny.html

[24] https://en.wikipedia.org/wiki/Tsiolkovsky_rocket_equation