

## Author's preface

Much of the story of mankind's first voyage to the moon centers on its engineering achievements. The sheer physical size of the vehicles and facilities not only dwarfed the previous generation of manned spacecraft, but still represents one of the greatest engineering achievements in history. But size alone is hardly the most valid metric for such an endeavor. Almost every conceivable figure of merit, whether pressures, temperatures, flow rates or stresses (and the list goes on) pushed the limits of the possible. But the magnitude of raw power overshadows equally impressive engineering at the other end of the spectrum. This book focuses on very different technologies, each with a unique set of attributes: fantastically small, breathtakingly precise, and in the most impressive feat of all, completely invisible to the eye. The Saturn V, with its 3,300 tons of structure and propellant would be helpless without a brain to give it a purpose. That brain, in the form of a small computer, guided the spacecraft and its precious human cargo to a new world, and home again.

In the 1960s, the Apollo Guidance Computer defined the state of the art in digital electronics, reliability and real-time computer control. However, the journey from a laboratory bench to a flight-ready system can only be described as tortuous, as the evolving requirements for a lunar mission forced nearly continuous redesigns. With so much time and effort already invested in an existing design, starting over with a new "clean sheet of paper" architecture was never an option. As a result, the AGC grew like a newlyweds' house. Over time, the happy couple's family grows, and so too must their house. Never willing to tear it down and begin anew, the house becomes a collection of architectural contradictions, with a blind hallway or two and plumbing that makes little sense. Still, the house is quite functional, with each iteration representing the best ideas at the time.

Some aspects of the AGC were truly revolutionary for their time. The exclusive use of integrated circuits in the processor ushered in a new era of computing, the novel memory design stored large amounts of data in a small space, and the human interface allowed real-time interaction with software. Other characteristics were like the newlyweds' house; the accumulated changes made the design awkward and inefficient. Finally, the severe lack of memory perpetually challenged the developers during the entire life of the project.

The *Apollo Guidance Computer: Architecture and Operation* presents the three major organizational levels of the AGC. First, the programmers' view of the AGC, known as the hardware architecture, is described in detail. This section discusses the key processor components and how they interact with each other. Importantly, the balance between power, flexibility and the limitations of the hardware is made evident by the designers' struggle to extend the capabilities of the system. The final design is full of exceptions and special cases, but is especially notable for its elegant solutions for optimizing every bit in memory. Understanding these constraints sheds much light on the decisions affecting the other parts of the system, from the Executive software to the operating procedures.

Efficiently running all the programs that are necessary for a complex spacecraft like Apollo demands a sophisticated operating system. The second section is devoted to the operating system software, known as the Executive, and a powerful execution environment called the Interpreter. All of the computer resources used by the mission programs were managed by the Executive. The Executive oversaw all of the processing necessary for an interactive, real-time computing environment: namely scheduling mission programs for execution,<sup>1</sup> accepting commands from the crew and displaying its results, and recovering from any errors that occur. Complex programming logic used in the mission programs required equally sophisticated capabilities that were unavailable in the hardware. A novel and compact language was designed which, by reflecting the problem to be solved, greatly reduced the workload on the programmer. This language was processed by a program called the Interpreter, which created the illusion of an entirely new and highly capable computer architecture. Slow, but rich with features, the Interpreter simplified the programming effort needed for the mission programs and greatly reduced their memory requirements.

The third and final section leaves the hardware and Executive behind to enter the world of mission software and its operation. Each mission phase, whether leaving Earth, landing on the Moon or performing a navigating sighting requires a variety of programs in the AGC. In turn, the execution of these programs must be coordinated with the operation of other spacecraft systems to satisfy a mission requirement. Properly covering each mission phase requires placing the AGC's operation into a larger perspective, by including the operation of selected spacecraft systems in the discussion. Each Apollo mission built on the experiences of its predecessors, so there is no "typical" mission to use as an example. Procedures and techniques reached their highest level of refinement in the latter missions – Apollo 15, 16 and 17 – so many of the descriptions will reflect these flights.

It is also worth noting what is *not* included in this book. With few exceptions,

<sup>1</sup> When discussing the wide variety of software used in computer systems, it is useful to separate their function into two distinct areas: "Executive software" or "operating system software", manages the resources and services of the computer. "Applications software" or "mission programs" focus on solving a particular problem for the user, whether it is calculating a payroll or the trajectory to the Moon.

there is no description of the underlying circuitry and chip technology, as the internal architecture of the AGC is invisible to both the programmer and the astronauts. To discuss both computer science and spaceflight in a single volume is a bold step in itself; to expand it to include electrical engineering would run the danger of overwhelming the reader. Also, despite its extensive treatment of the AGC hardware, Executive and Interpreter, this book is not intended to be an AGC programming tutorial. Including material on programming techniques, especially necessary for novice readers, followed by the obligatory problem exercises would move this book far from its intended scope. Finally, descriptions of the architectural evolution of the hardware and software are consciously avoided. Specifically, the book reflects the Block II configuration of the AGC, which was used in all manned Apollo flights. The compressed timeline to achieve a manned lunar landing meant there were several concurrent development projects at any one time, and a description of every variant would result in a hopelessly complex account.

*Apollo Guidance Computer: Architecture and Operation* is intended to be the essential reference for both spaceflight historians and those with computer or engineering backgrounds. Bringing those disciplines together into a single work is not just an interesting academic exercise; it offers a new way to understand the mission design process. Both the design decisions and the compromises made during the AGC's development directly influenced the overall operation of the Apollo spacecraft. Once defined, the spacecraft's operational capabilities became the principal driver in creating mission objectives and procedures. By understanding the AGC's capabilities and constraints, the reader will gain important insights into how the lunar missions assumed their final form.



<http://www.springer.com/978-1-4419-0876-6>

The Apollo Guidance Computer  
Architecture and Operation

O'Brien, F.

2010, XVIII, 440 p. 30 illus., Softcover

ISBN: 978-1-4419-0876-6

A product of Praxis