

Introduction and History

What an operating system does

Without an operating system (OS), a computer would be just a really complex electronic circuit that is extremely tedious to use. Here are a few highlights of what operating systems do:

1. They make it so that there can be multiple programs running at once, and the ones that need to use the CPU right away can get priority.
2. They manage memory so that each program's memory is protected from being read or corrupted by other programs.
3. They abstract away the nitty-gritty details of hardware so that people and programmers can just click the on "save" or call the save() function without having to know all the 1000 pages of details that comprise a hard drive's user manual.

Computing in the 1940s

The first digital computer was the Atanasoff-Berry computer from the early 1940s. Other computers started to be built in the 1940s, the most famous of which was the ENIAC. These computers were very large, in some cases taking up an entire large room. Some were programmed by an operator that moved around wires on a plugboard. Programs were mostly written in pure numerical machine language. By the late 1940s, assembly language had been developed which gave human readable names to the numerical machine language instructions. Operating systems at this point were essentially nonexistent. Most computing was for scientific and military calculations.

1950s

In the 1950s, instead of a single operator running programs on a computer, a new system developed. You would write your program out by hand and transfer it to punched cards. These cards were fed into a machine that wrote the program onto a tape. That tape was then fed into the computer, which would then output the result, often on another tape. By the 1950s computers started to be used by large businesses to process data, especially in banking and insurance. The first programming languages were developed. Some of those early languages, including Fortran, Cobol, and LISP, are still in use today.

Computer time was precious, and a lot of time was wasted loading and unloading tapes, so a new system called *batch processing* was developed. In it, several programs were all collected into a single batch on the same tape, and the computer would run them one after another. The program that ran the batch was essentially a very early operating system. A downside for programmers was that they would have to wait hours between writing their program and seeing its result. Batch processing is still used, though not with punched cards anymore. For example, some banks group transactions into batches that are run maybe once a day after business hours.

1960s

In the 1960s computers began to be built with transistors instead of vacuum tubes, which made them much smaller and more reliable. Operating systems began to get more sophisticated. Two important operating systems developed in the 1960s are OS/360 and Multics.

OS/360 was developed by IBM in the mid to late 1960s. It was originally supposed to be an all-encompassing operating system that could run on a wide range of machines. The project itself turned into something of a boondoggle, costing hundreds of millions of dollars, being well behind schedule, and having many bugs. Many

good software engineering lessons were learned from that project. OS/360 was a batch processing system whose descendants are still in use today on large mainframe computers.

One of its most important contributions is the idea of *multiprogramming*, where multiple programs could be active at the same time. One early type of multiprogramming is as follows: programs often spend a lot of their time doing input/output (I/O) operations. These are things like writing to a printer, reading a file from disk, waiting for someone to type a key, etc. These operations don't require the CPU very much, and so if the program that currently has control of the CPU starts doing I/O, the operating system can bring in another program to run while the original is busy with its I/O.

Multics was developed by Bell Labs, GE, and MIT in the late 1960s. The first two dropped out, and MIT eventually released it, though just like with OS/360, the project was a little too ambitious and it ended taking much longer to develop than expected. Multics had a small, but enthusiastic user-base. It is best known now for introducing some key ideas, like virtual memory and hierarchical file systems, that came to be used in later operating systems.

Whereas OS/360 was a batch processing system, Multics was a *time-sharing* system. In batch processing, there is not really much user interaction. You load up the batch, it runs, and you get the output. Time-sharing systems worked differently. Here you had the main computer and various terminals, which might not be much more than a keyboard and monitor, a little like today's thin clients. Each user at a terminal shared the main computer with others, with the main computer's OS giving each user at a terminal a little bit of CPU time before rotating on to the next user.

1970s

In the early 1970s, Ken Thompson of Bell Labs, who had worked on the Multics project, decided to make a scaled-down version, which he called Unics and eventually Unix. It became pretty popular at Bell Labs, and more things were added to it. Dennis Ritchie at Bell Labs developed the C programming language in order to write tools for Unix, and it was also used to rewrite the OS itself. Unix was shared freely with other institutions and it quickly became popular.

By the mid 1970s, a few companies, particularly Intel and Motorola, were producing *microprocessors*. These made it possible to build computers small enough and cheap enough that ordinary people could have them. The earliest successful one was the Altair 8800, which was actually a kit for electronics hobbyists. In the late 1970s, Apple was formed and started selling computers with keyboards and monitors to home users.

1980s and 1990s

In the early 1980s, IBM got into the personal computer business. Intel-based computers in the 1970s primarily ran an operating system called CP/M. When IBM decided to introduce their PC, they initially tried to use CP/M, but due to a misunderstanding, it didn't work out. Microsoft then agreed to put together an operating system for IBM PCs, even though Microsoft didn't yet do operating systems. They bought a company that produced a clone of CP/M called the quick and dirty operating system (QDOS). Microsoft reworked it into the Microsoft Disk Operating System (MS-DOS), which became the main operating system for PCs for many years.

In the 1970s, Xerox, mostly known for the copy machine business, had established a research center in California, called Xerox PARC. They hired a bunch of really smart people and let them work. They developed many of the things we associate with modern computing, such as graphical user interfaces (GUIs) with moving windows and bitmapped graphics, laser printing, and ethernet. Xerox, however, did not take advantage of the developments at their research lab at the time. Instead, it was Apple. Xerox had invested some money in Apple, and as part of the deal, people from Apple were allowed to visit Xerox PARC. They took some of the ideas and used them to eventually develop the Apple Macintosh, which came out in 1984. Microsoft soon copied many of those ideas to produce their own graphical OS, Windows. Windows wasn't really successful until version 3 came out in 1990. Those early versions of Windows were really just graphical shells running on top of the non-graphical MS-DOS operating system. Later versions of Windows, starting with Windows 95 were truly GUI-based.

In the mid 1980s at UC Berkeley, a version of Unix called BSD was developed. This is still in use today. BSD was adopted and modified by the computer company NeXT for their computers' operating system. The company was acquired by Apple in the late 1990s, and that operating system became the basis of MacOS X and later iOS.

In the late 1980s, Computer Science professor Andrew Tanenbaum developed a small version of Unix, called Minix, for use in teaching classes in operating systems. A young Finnish programmer, Linus Torvalds, took off on the Minix idea and created his own version, Linux, as an open source program. Linux has since grown to be used for many purposes. The Android OS is a modified version of Linux, which makes Linux the most widely used operating system in the world.