

A History of the GUI

By [Jeremy Reimer](#)

Thursday, May 05, 2005

[arstechnica.com](#)

Contents

Introduction	1
Prehistory	1
The father of the GUI	2
The mother of all demos	3
Xerox Park	5
Smalltalk	8
Apple	10
Other GUIs during the 1980s	12
VisiOn	12
Windows 1.0	13
Tandy DeskMate	15
GEM	15
Amiga Workbench	16
GEOS	17
Windows 2.0	17
Acorn	17
NeXTSTEP	18
OS/2	19
X marks the spot	20
The 1990s and beyond	21
Windows	22
OS/2 in the 1990s	23
BeOS	24
Mac OS X and Aqua	25
Conclusions	26
GUI development timeline	28

Introduction

Today, almost everybody in the developed world interacts with personal computers in some form or another. We use them at home and at work, for entertainment, information, and as tools to leverage our knowledge and intelligence. It is pretty much assumed whenever anyone sits down to use a personal computer that it will operate with a graphical user interface. We expect to interact with it primarily using a mouse, launch programs by clicking on icons, and manipulate various windows on the screen using graphical controls. But this was not always the case. Why did computers come to adopt the GUI as their primary mode of interaction, and how did the GUI evolve to be the way it is today?

Prehistory

Like many developments in the history of computing, some of the ideas for a GUI computer were thought of long before the technology was even available to build such a machine. One of the first people to express these ideas was Vannevar Bush. In the early 1930s he first wrote of a device he called the "Memex," which he envisioned as looking like a desk with two touch screen graphical displays, a keyboard, and a scanner attached to it. It would allow the user to access all human knowledge using connections very similar to how hyperlinks work. At this point, the digital computer had not been invented, so there was no way for such a device to actually work, and Bush's ideas were not widely read or discussed at that time.

However, starting in about 1937 several groups around the world started constructing digital computers. World War II provided much of the motivation and funding to produce programmable calculating machines, for everything from calculating artillery firing tables to cracking the enemy's secret codes. The perfection and commercial production of vacuum tubes provided the fast switching mechanisms these computers needed to be useful. In 1945,

Bush revisited his older ideas in an article entitled “As We May Think,” which was published in the Atlantic Monthly, and it was this essay that inspired a young Douglas Engelbart to try and actually build such a machine.

The father of the GUI



Douglas Engelbart in 1968

Douglas Engelbart completed his degree in electrical engineering in 1948 and settled down in a nice job at the NACA Institute (the forerunner of NASA). However, one day while driving to work he had an epiphany: he realized that his real calling as an engineer was not to work on small projects that might only benefit a few people. Instead, he wanted to work on something that would benefit all of humanity. He recalled Bush’s essay and started thinking about ways in which a machine could be built that would augment human intellect. During the war he had worked as a radar operator, so he was able to envision a display system built around cathode ray tubes where the user could build models of information graphically and jump around dynamically to whatever interested them.

Finding someone to fund his wild ideas proved to be a long and difficult task. He received his PhD in 1955, and got a job at the Stanford Research Institute, where he received many patents for miniaturizing computer components. By 1959 he had earned enough recognition to receive funding from the United States Air Force to work on his ideas. In 1962, Douglas published his ideas in a seminal essay entitled “Augmenting Human Intellect.” In this paper, Douglas argued that digital computers could provide the quickest method to “increase the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems.” He envisioned the computer not as a replacement for human intellect, but a tool for enhancing it. One of the first hypothetical examples he described for this technology was of an architect designing a building using something similar to modern graphical CAD software.

This was a huge leap in thinking for 1962. The only computers that existed at the time were giant mainframes, and typically users would interact with them using what was called “batch processing.” A user would submit a program on a series of punch cards, the computer would run the program at some scheduled time, and then the results would be picked up hours or even days later. Even the idea of having users enter commands on a text-based terminal in real-time (called “time-sharing” in the jargon of the day) was considered radical back then.

Douglas and his growing staff worked for years to develop the ideas and technology that finally culminated in a public demonstration in front of over a thousand computer professionals in 1968.

The mother of all demos

Englebart’s demonstration was a multimedia extravaganza: television cameras were trained on his face, his hands, and the small display screens he was viewing. The cameras faded in and out between all three views, often showing more than one at the same time. This level of demonstration was necessary because so many of the concepts being shown were completely new. The system was called NLS, or oN-Line System, because it was also networked between multiple computers.

The display system was based on vector graphics technology and could display both text and solid lines on the same screen. Because of limited memory space in the mainframe computer, it could only display upper-case characters, although true upper-case was displayed by the use of a short horizontal line directly above any capitalized letters.



The oN-Line System display, keyboard and mouse

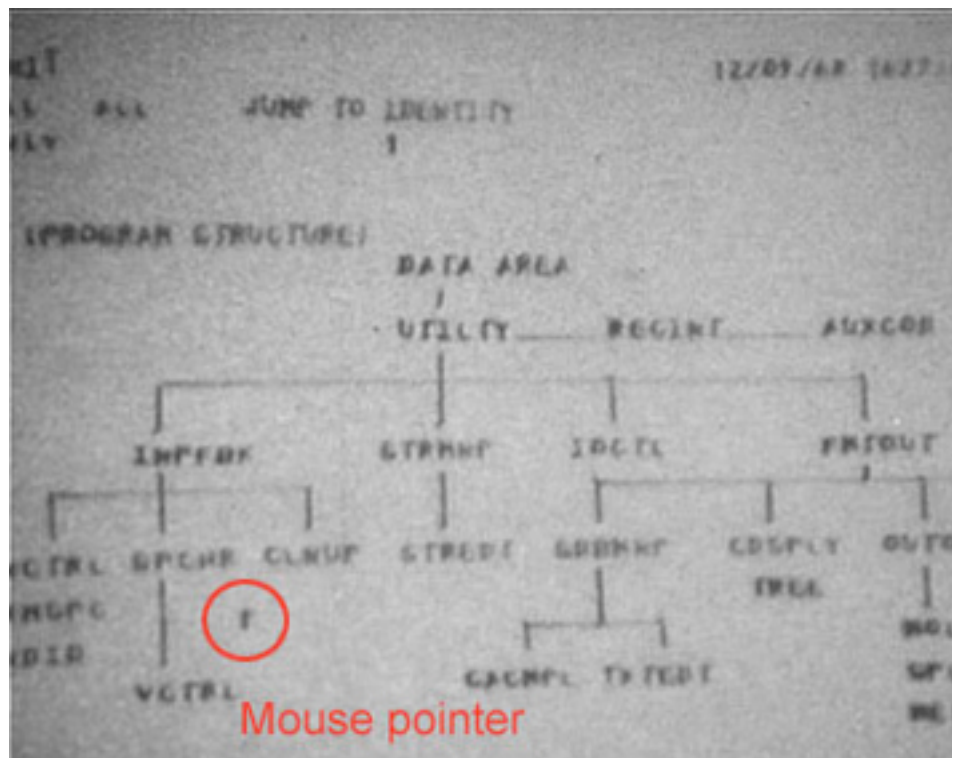
Douglas' hands operated three input devices: a standard typewriter-style keyboard, a five-key "chording keyboard" (combinations of the five keys could produce 2^5 or 32 separate inputs, enough for all the letters of the alphabet), and a small rectangular box about the size of a couple of juice boxes with three buttons near the top, connected to the computer with a long wire. This was the mouse, invented by Douglas himself and built by one of his engineers. Nobody knew who first started calling it a mouse, but the name stuck back then, and has remained ever since. Mechanically it was slightly different from modern mice in that the two circular wheels connected to the internal potentiometers rolled directly on the table surface, instead of being manipulated by a single mouse ball rubbing against rollers. However, to the end user it operated virtually identically to a modern mouse. Other input devices had been tried (such as touch screens and light pens), but user testing found the mouse to be the most natural way to manipulate an on-screen cursor. This remains true today.



Close-up of the keyboard and mouse

With the invention of the mouse came the invention of the mouse pointer, which in this system was a stick arrow, about the height of a single character, pointing straight up. This was called a "bug" by Douglas' team, but this term did not survive into modern use. When objects were selected, the "bug" would leave dots on the screen to mark this action.

Many of the things demonstrated in this marathon session seemed to come from decades in the future, and most of the people watching it were hard pressed to even understand all of what they were seeing. The demo featured hypertext linking, full-screen document editing, context-sensitive help, networked document collaboration, e-mail, instant messaging, even video conferencing! It was all somewhat overwhelming, and due to limitations in the video system, sometimes difficult to tell what was going on. For example, NLS supported multiple windows, but there was no obvious way to indicate boundaries between them (such as window borders, title bars, etc)



A screen from the NLS demo. The red circle indicates the mouse pointer.

Douglas and his staff would continue working on these revolutionary ideas until the institute was closed down in 1989 due to lack of funding. However, there was never enough money to turn these ideas into commercial products. That job was undertaken by a company that most people associate with old-fashioned paper.

Xerox Park

Douglas Englebart's demonstration in 1968 amazed many people and overwhelmed many more. It opened people's eyes to what could be possible in the future — a future where people all over the world collaborated on electronic documents displayed on computer screens and transmitted instantly over networks to other users. Such a future did not bode well for a company that had made its fortune selling photocopiers.

Xerox upper management, fearing the inevitable demise of their paper-based company in the "paperless" future, decided that they had better make sure they controlled this new technology. They formed the Palo Alto Research Center, or PARC, in 1970.

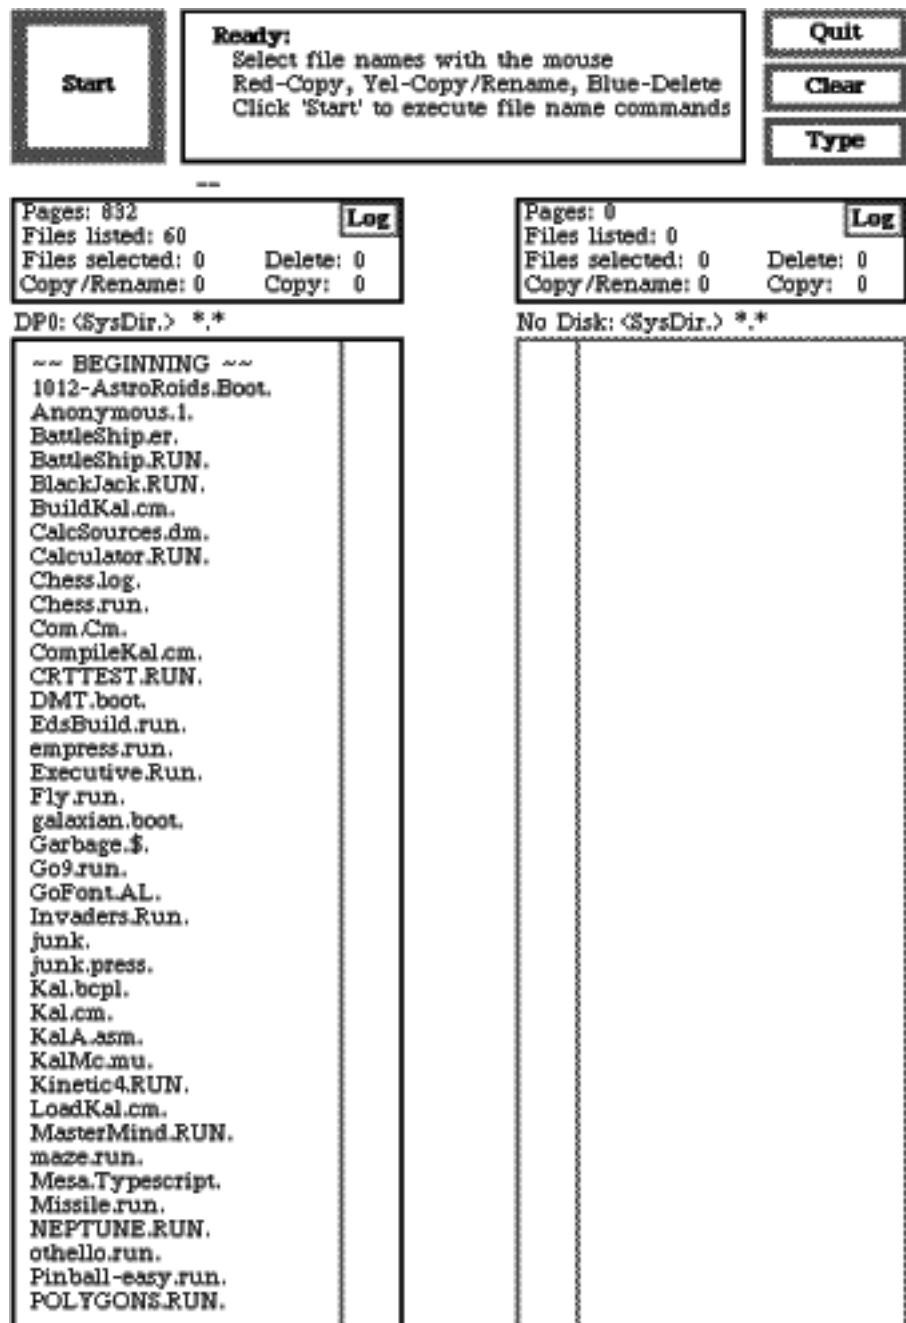


The Xerox Alto

People came to PARC to work on five-year projects that were their dreams, and were given great freedom and latitude to do whatever interested them. This allowed the group to attract the top computer science researchers in the country. The atmosphere was relaxed, yet charged with excitement. PARC researchers truly believed they were inventing the future of computing, and they ended up doing just that.

One of the first things they invented was the laser printer, a natural complement to Xerox's copier business. But such a printer demanded a more graphical way for a computer to prepare documents to begin with. Since there were no such computers at the time, PARC invented their own. This was dubbed the Alto, and was first completed in 1973.

The Alto was not a microcomputer as such, although its working components did fit in a minibar-sized tower that fit under the desk. Its most striking feature was its display, which was the same size and orientation as a printed page, and featured full raster-based, bitmapped graphics at a resolution of 606 by 808. Each pixel could be turned on and off independently, unlike typical terminals of the time which could only display fixed text characters, and unlike the vector-based NLS which could only display text and straight lines. It also had a keyboard and a modernized version of Englebart's mouse, again with three buttons. The mouse cursor itself became a bitmapped image, and for the first time took the familiar diagonal-pointing arrow shape we know today, as well as morphing into other shapes depending on the task being performed.



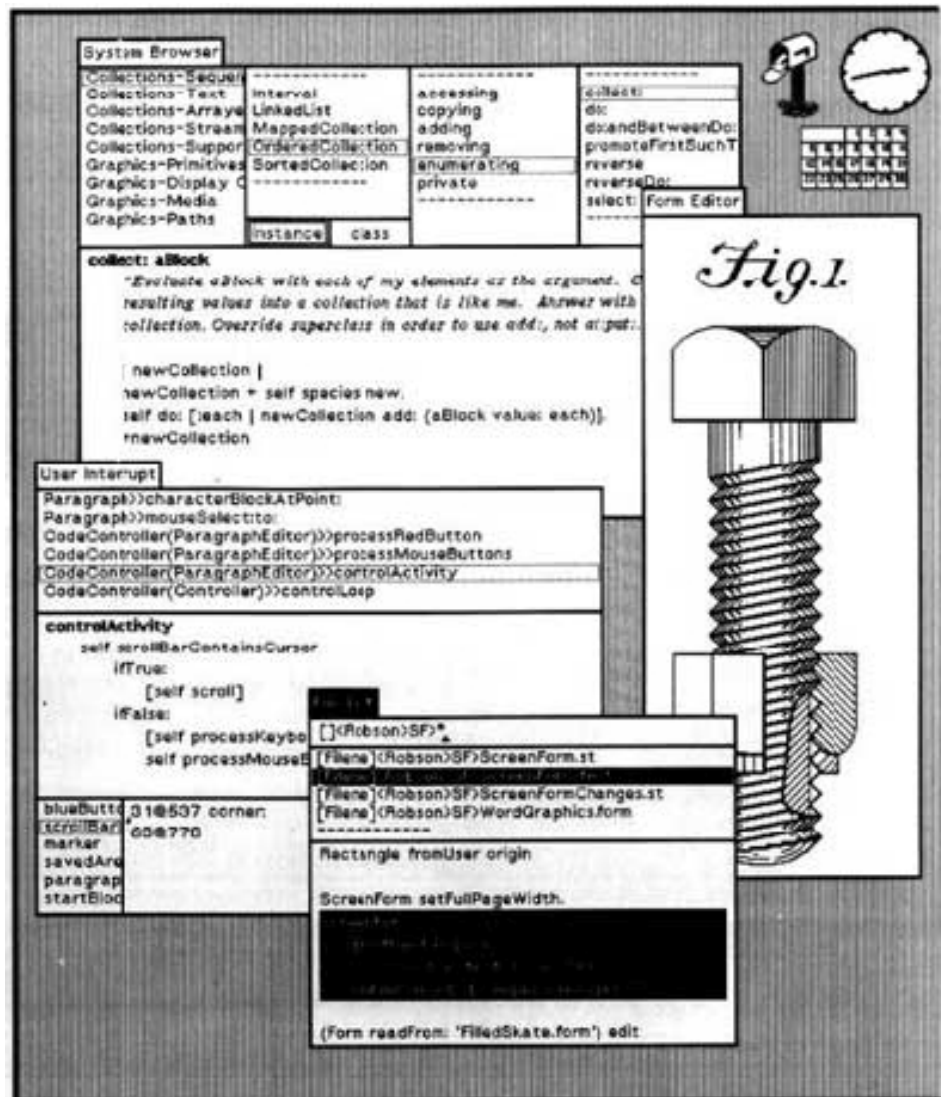
The Alto File Manager

The first software written for the Alto was rather crude, and only slightly graphical. For example, the file manager displayed directory listings in two columns (like Norton Commander) surrounded by boxes, but there were no “windows” as we know them today. A graphical word processor, called Bravo, was developed that could display different fonts and text sizes on the screen at the same time, but it had a slightly different user interface with menus on the bottom rather than on the top. (The author of Bravo, Charles Simonyi, would eventually join Microsoft and recreate his work as the original Word for DOS.) There was a bitmapped graphics editor that worked a lot like Paint does today, but it too had its own, different user interface. The PARC researchers realized that what was needed was a consistent user interface for new applications, and to make that happen a whole new visual code development environment would have to be invented. This was Smalltalk, the first modern GUI.

Smalltalk

Smalltalk was conceived as a programming language and development environment so easy to use that a child could understand it, and in many respects was successful in this goal. Smalltalk was the world's first object-oriented programming language, where program code and data could be encapsulated into single units called objects that could then be reused by other programs without having to know the details of the object's implementation. It also had modern, Java-like features like automatic memory management, to take some of the hard work away from the programmer. The development environment of Smalltalk was also the user interface that Smalltalk programs ran in, and introduced many modern GUI concepts. It first began to take shape around 1974, and was continuously updated and enhanced.

SmallTalk was a graphical development environment (like a modern visual IDE), but it was also the environment that it ran under. It was kind of like if Microsoft had developed Visual Studio as a single application that was itself Windows. You would launch Smalltalk from the file manager just like a regular application, but once it launched it took over the whole way the Alto's environment was presented.




The SmallTalk development and GUI environment

Individual windows in Smalltalk were contained by a graphical border, and stood out against the grey pattern of the background below them. They each had a title bar on the top line of each window which could be used to identify the window and move it around the screen. Similar to BeOS, the title bar did not stretch the full length of the window, but started at the top left and only extended as far as the title itself. Windows could overlap other windows on the screen, and a selected window would move itself to the top of the “stack.” The concept of “icons” was also invented at this time — small iconic representations of programs or documents that could be clicked on to run them or manipulate them. Popup menus were also invented at the same time — the user would click one of the mouse buttons and hierarchical, graphical menus based on the task at hand would appear at the last position of the mouse cursor. Also appearing for the first time were scroll bars, radio buttons and dialog boxes.

The combination of Smalltalk and the Alto was essentially a modern personal computer with a very similar graphical user interface to the ones we use today. Altos had networking and could send e-mail to and receive it from one another, and seemed ideal for an office environment. Many of the PARC team wanted Xerox to market the new, cost-reduced Alto III as a commercial product (the original Alto was never available for sale) but Xerox management declined.

Eventually, a stripped-down version of the Alto, the Xerox Star 8010 Document Processor, was released to the public in 1981 for US \$17,000. The Star had some differences from the Alto, most significantly the ability to overlap windows was removed as it was thought too confusing for the general public. Instead, the Star used tiled windows. As significant as the Star’s release was, it was too little and too late for Xerox, who had by that time lost most of its top researchers to other companies.

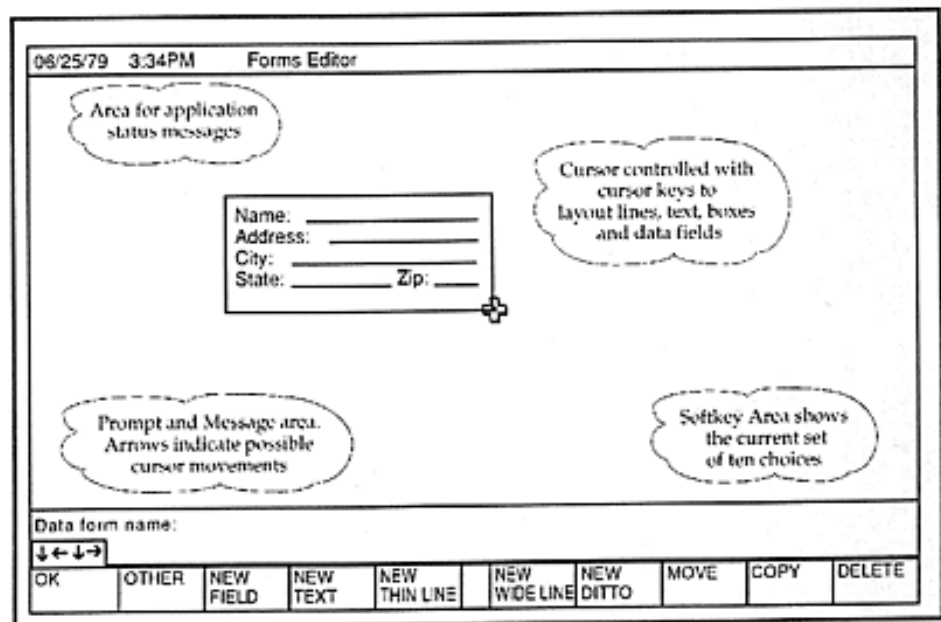


Screenshot from the Xerox Star. Note the tiled application windows.

Apple

The most important of these GUI pioneers was a small startup founded in a garage in 1976 by Steve Jobs and Steve Wozniak, called Apple Computer. Apple had built its fortune on the wildly popular Apple II, which displayed both text and graphics but had a traditional command line interface. Apple was a young company that found itself flush with money, and was more willing to take risks. Many former Xerox PARC engineers found new jobs with Apple, and set about to recreate their work on the Alto and Smalltalk but on a product that would actually see commercial release and potentially become very popular.

Work on Apple's next-generation Lisa computer, which had started life as a traditional text-based command line computer for business use, was transformed by the influx of PARC people. Steve Jobs himself became a convert to the GUI religion when his employees arranged a grand tour for him of the PARC facility. The Lisa would henceforth become a graphical computer, but the design of the interface was not yet set in stone.



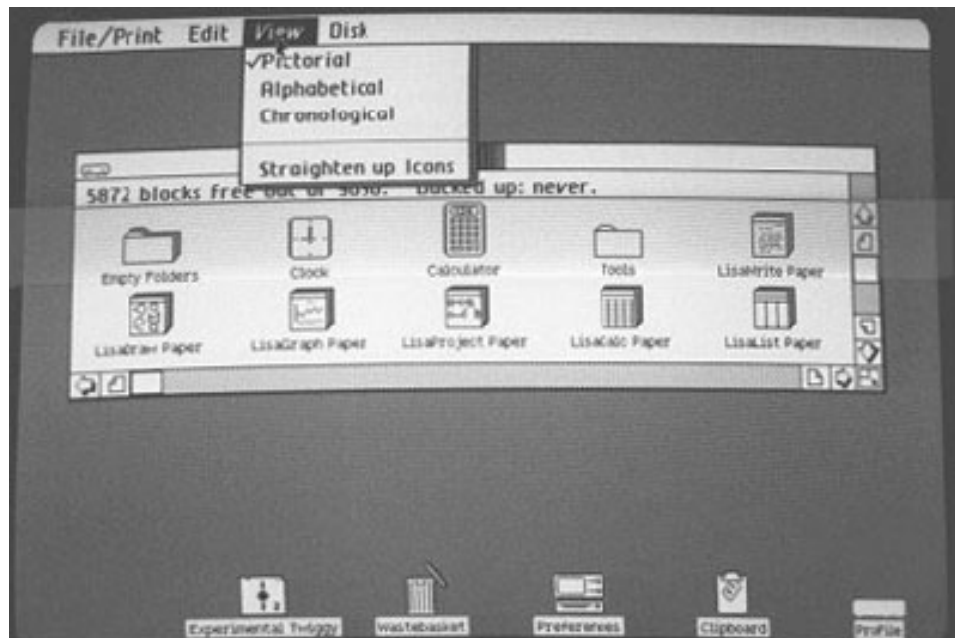
Early Lisa user interface mockup, circa 1979.

Many different prototypes for the Lisa interface were mocked up on the Apple II, including a task-based interface dubbed "Twenty Questions" because it seemed to take that long to get the computer to do anything, and a four-column file browser similar to one that appeared with Smalltalk and would reappear much later in NeXTSTEP and Mac OS X. The Lisa team eventually settled on an icon-based interface where each icon indicated a document or an application, and developed the first pull-down menu bar, where all menus appeared at the very top line of the screen.

Other innovations from the Lisa team included the idea of checkmarks appearing next to selected menu items, and the concept of keyboard shortcuts for the most frequently used menu commands. The Lisa also changed some PARC conventions, such as eschewing proportionally-sized scroll bars for fixed-height ones, and added new conventions, such as a trash can for dragging documents scheduled for deletion, and the idea of "graying out" menu

options if they were not currently available. The three-button mouse, which had been changed to a two-button design in the Star for simplicity, was further simplified to have only one button for the Lisa. As the interface required at least two actions for each icon (selecting and running) the concept of double-clicking was invented to provide this functionality. Double-clicking would later become a standardized way for all GUIs to launch a program, even those with multiple-button mice.

The Lisa user interface invented other GUI concepts that we still use today. While SmallTalk and the Xerox Star had icons that represented files, the Lisa interface was the first to have the idea that icons could represent all files in the file system, which could then be browsed through using a hierarchical directory structure where each directory opened in a new window. The idea of “drag-and-drop” was also invented at this time, and the concept of using drag and drop to do file manipulation (for example, selecting a group of files with the mouse and then dragging them to a new folder to copy them) naturally developed from this concept. Less visible but still important to the user interface was the idea of “resource forks,” which embedded information about a file separately from the file itself, and the idea of “creator classes” meant that each file could be assigned an application that would launch when you double-clicked on that file.

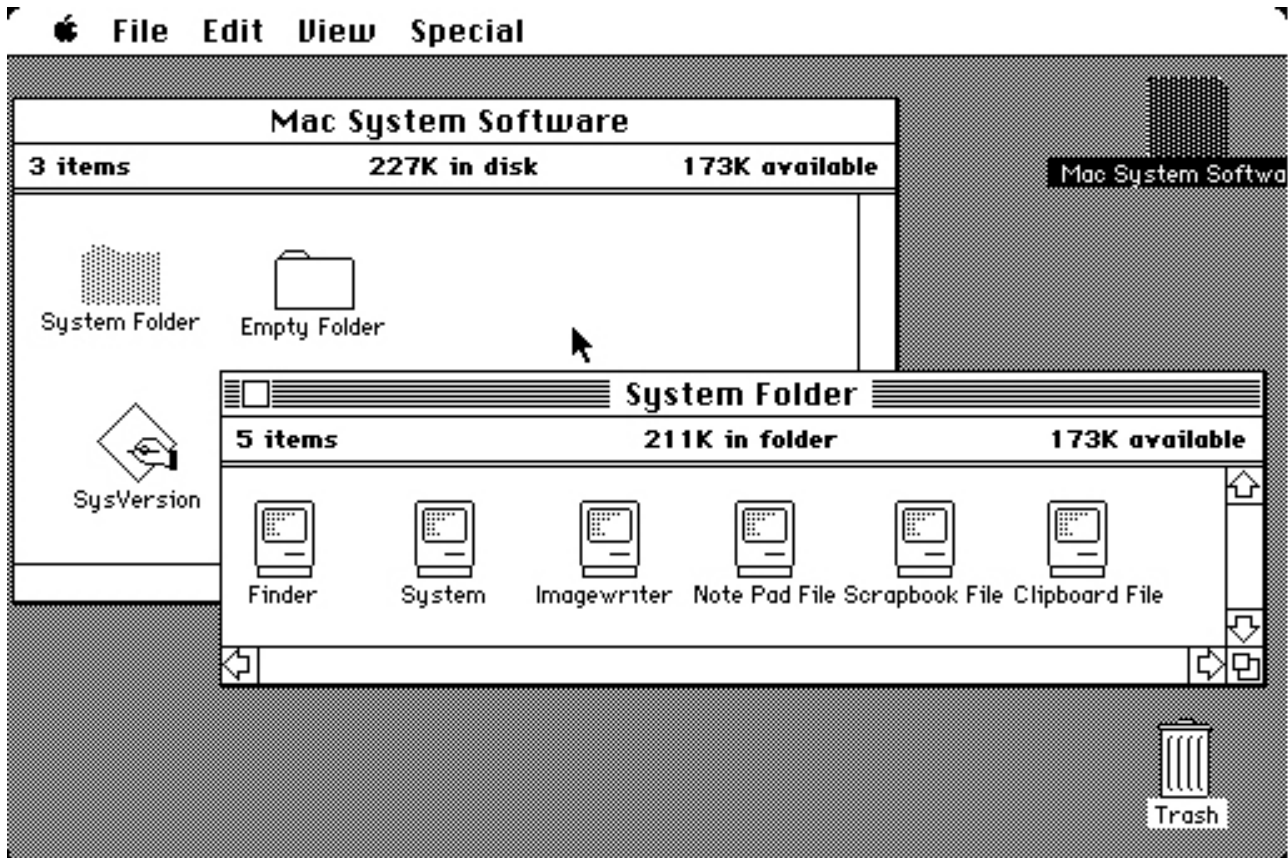


Final Lisa user interface.

One critical advance from the Lisa team came from an Apple engineer who was not a former PARC employee, but had seen the demonstration of Smalltalk. He thought he had witnessed the Alto's ability to redraw portions of obscured windows when a topmost window was moved: this was called “regions”. In fact, the Alto did not have this ability, but merely redrew the entire window when the user selected it. Despite the difficulty of this task, regions were implemented in the Lisa architecture and remain in GUIs to this day. (Try rapidly moving, say, an Explorer window over top of a Word document. If your eyes are very fast, you can still see which parts of the window are redrawn and which are not.)

Work on Lisa started in 1979 but the computer was not released until 1983. Despite its advanced features, the price tag of US \$10,000 and the difficulty of writing software for the new machine limited sales. A low-cost, stripped down

version of the Lisa was needed, and this task fell to Steve Jobs himself. His Macintosh project achieved the goal of a lower-cost graphical computer by shipping with a small 9-inch monochrome screen (512 by 384 pixels), a mere 128 kilobytes of memory, no multitasking or even the ability to task switch between more than one program, and a single floppy drive. This was introduced to the world in dramatic fashion in 1984, selling for US \$2,495. It retained most of the GUI features of the Lisa, and even shared some of its low-level code, but the operating software itself was written from scratch to fit in the small memory footprint.



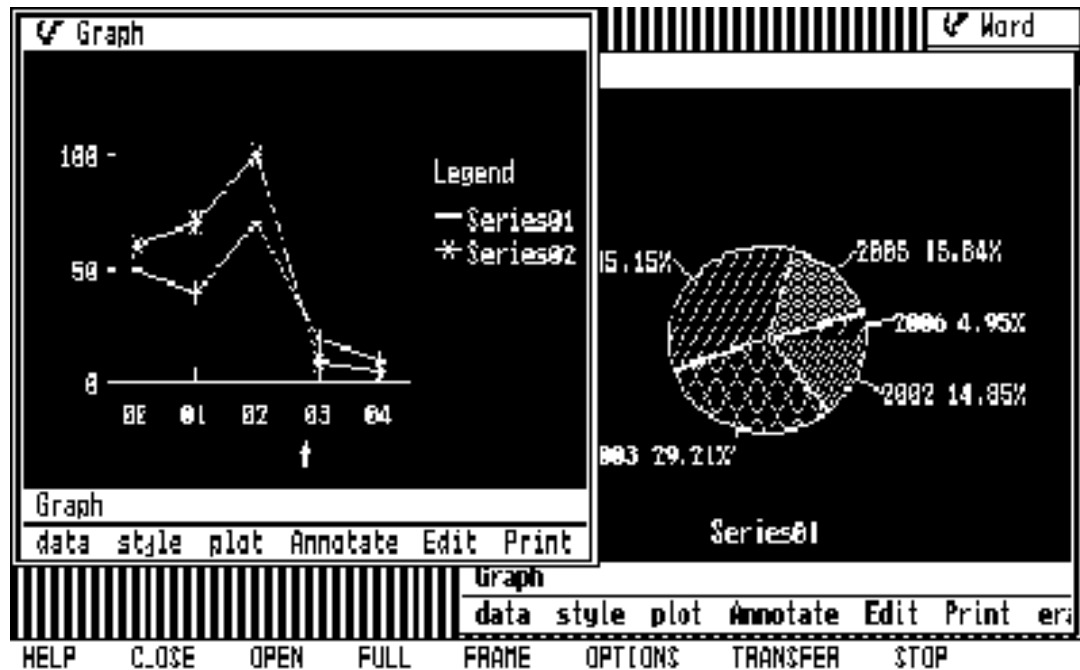
Macintosh System 1

Other GUIs during the 1980s

VisiOn

By this time more than one company besides Apple was working on graphical user interfaces for personal computers. VisiCorp, makers of the first ever spreadsheet VisiCalc, was working on a graphical interface for IBM PCs. It was

released as VisiOn in 1983, but the extremely high price (US \$1,495 for the software alone) and steep hardware requirements for the time (PC-XT with hard disk, 512KB of memory, and a mouse) kept it from being a big seller. In addition, it was fairly clunky, using the monochrome CGA graphics mode (640 by 200 pixels) and being much more text-based than other GUIs. For example, it did not use icons at all, requiring the user to click on text labels to start programs or work with documents. It did not use proportionally-spaced fonts, the way the Alto, Star, Lisa and Macintosh had; instead, all characters were of fixed width. It even abandoned the diagonally-pointing mouse cursor arrow, reverting to 1968 with a vertical one!

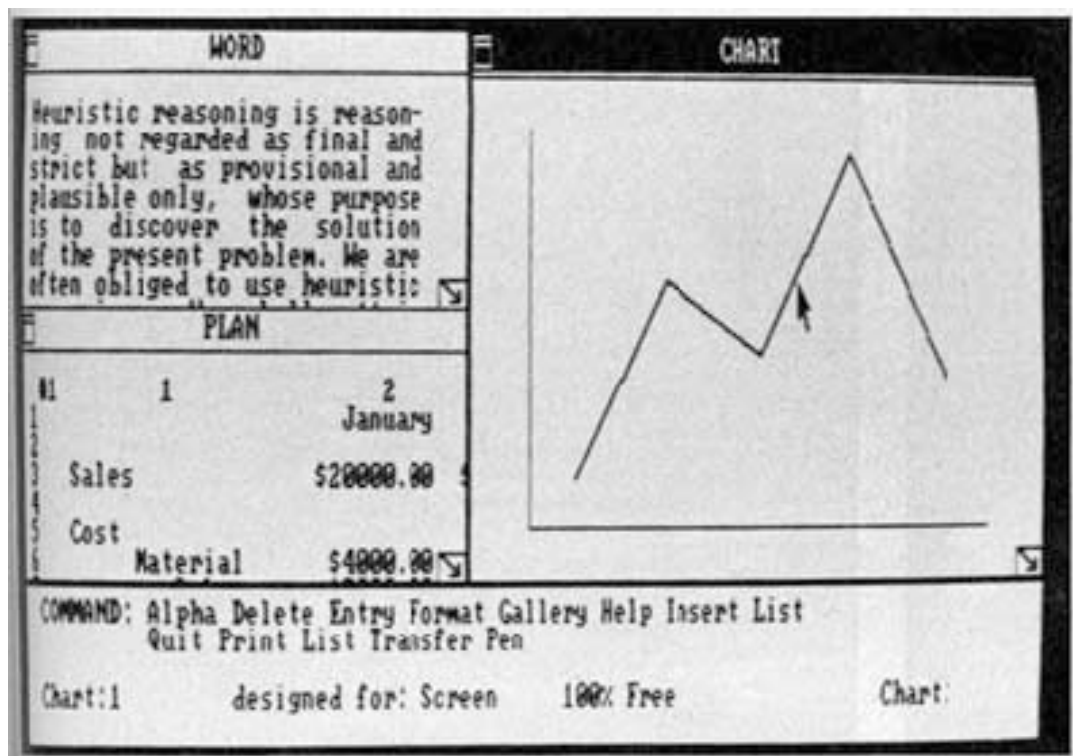


VisiOn user interface

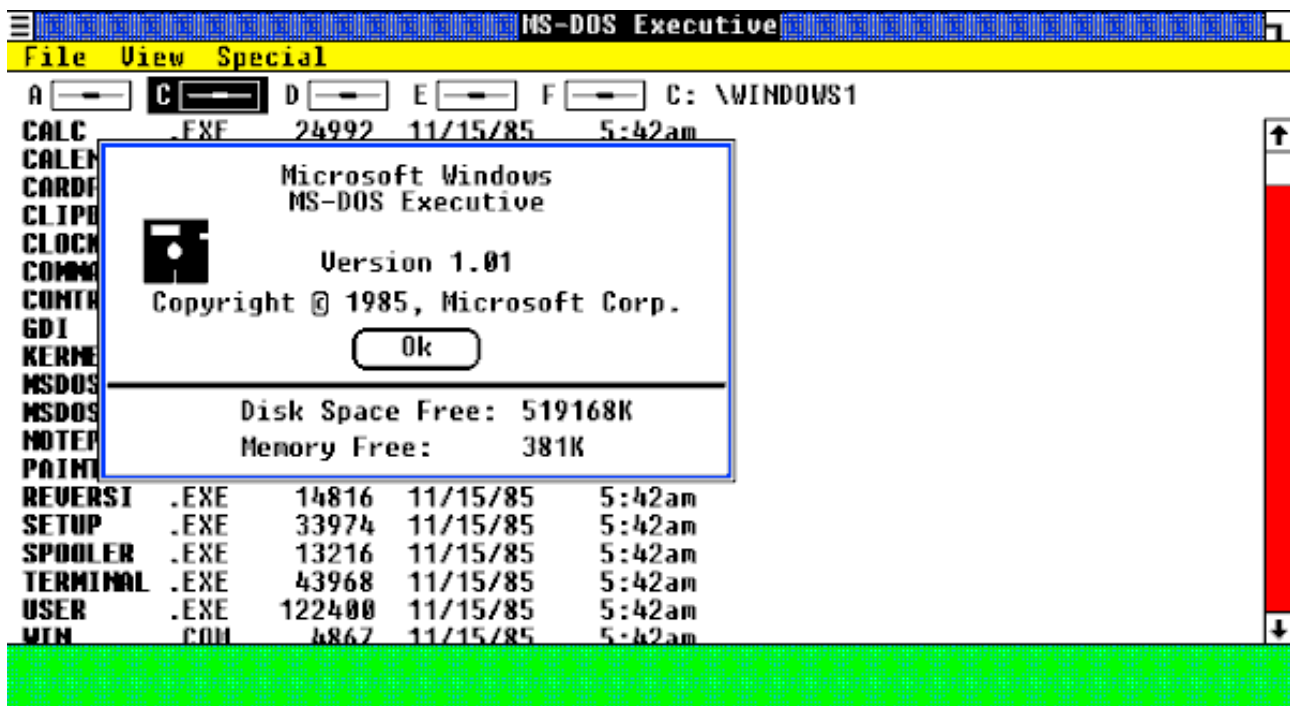
Windows 1.0

Overall, VisiOn was clunky and a failure, but its mere announcement inspired Bill Gates to come up with a competing product, initially titled Interface Manager, but later renamed to the somewhat more familiar Windows. Announced in 1983, early screenshots looked like a cross between VisiOn and the Microsoft Word for DOS interface (which itself descended from Bravo, the first GUI word processor on the Alto), but by the time it was released in 1985 it was in color and had all the usual GUI trappings, such as scrollbars, window control widgets, and menus, although instead of a single menu bar as on the Lisa and Macintosh, each application had its own menu bar attached to it, just below the title bar.

Another departure was the use of tiled, rather than overlapping windows. This decision was made by the same people from Xerox PARC who had decided on tiled windows for the Star to avoid confusing users, but Bill Gates didn't like it very much and all future versions of Windows would use the overlapping model. Microsoft was one of the earliest 3rd party developers for the Macintosh, and actually got to use beta models of the first Mac before it was released to the world. Undoubtedly this influenced the direction of future releases of Windows.



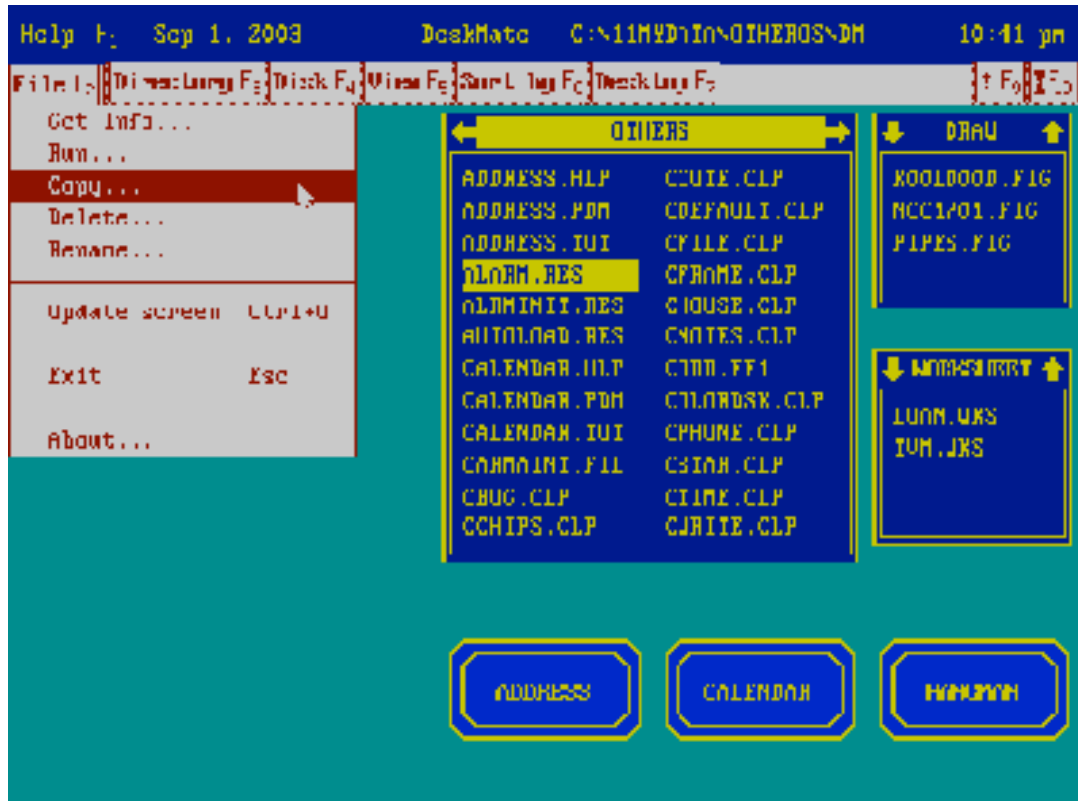
Windows prototype, circa 1983. Note the similarity to Word for DOS.



Windows 1.01

Tandy DeskMate

Tandy Computers released the first version of their own GUI in 1984. Called DeskMate, it was designed to be used primarily with the keyboard, using function keys as shortcuts for menus, and did not have overlapping windows. It was quite difficult to use and did not achieve much success beyond being bundled with Tandy PCs for a few years.



Tandy DeskMate.

GEM

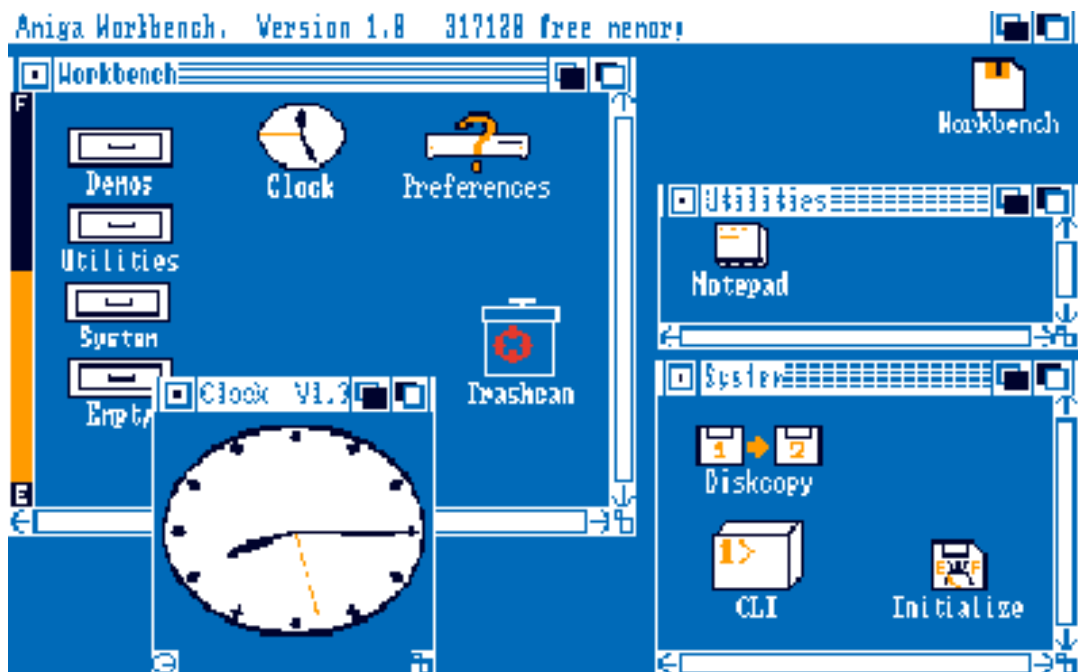
Later in 1985 Digital Research introduced GEM, a windowing GUI for DOS, and also for Atari's new ST computer that was released in the same year. GEM was quite similar to the Lisa/Macintosh GUI, so similar in fact that Apple would sue Digital Research, with the result that they ultimately castrated the PC version. Strangely, the Atari version was allowed to continue untouched. It also used a single menu bar at the top of the screen.



GEM 1.0 running on the Atari ST computer.

Amiga Workbench

Commodore's Amiga computer, introduced later that year, came with its own GUI, Workbench. It featured some new ideas, such as the ability to move windows up and down the "stack", and the ability to select, move, and work in a window without automatically bringing it to the front. It also had a single menu bar at the top that was normally hidden from view and activated using the right mouse button.



Workbench running on the Amiga 1000.

GEOS

In 1986 Berkeley Softworks released GEOS, a GEM-like GUI for older computers like the Commodore 64 and Apple II. It was later ported to the PC as GeoWorks and briefly became a competitor to Microsoft Windows.



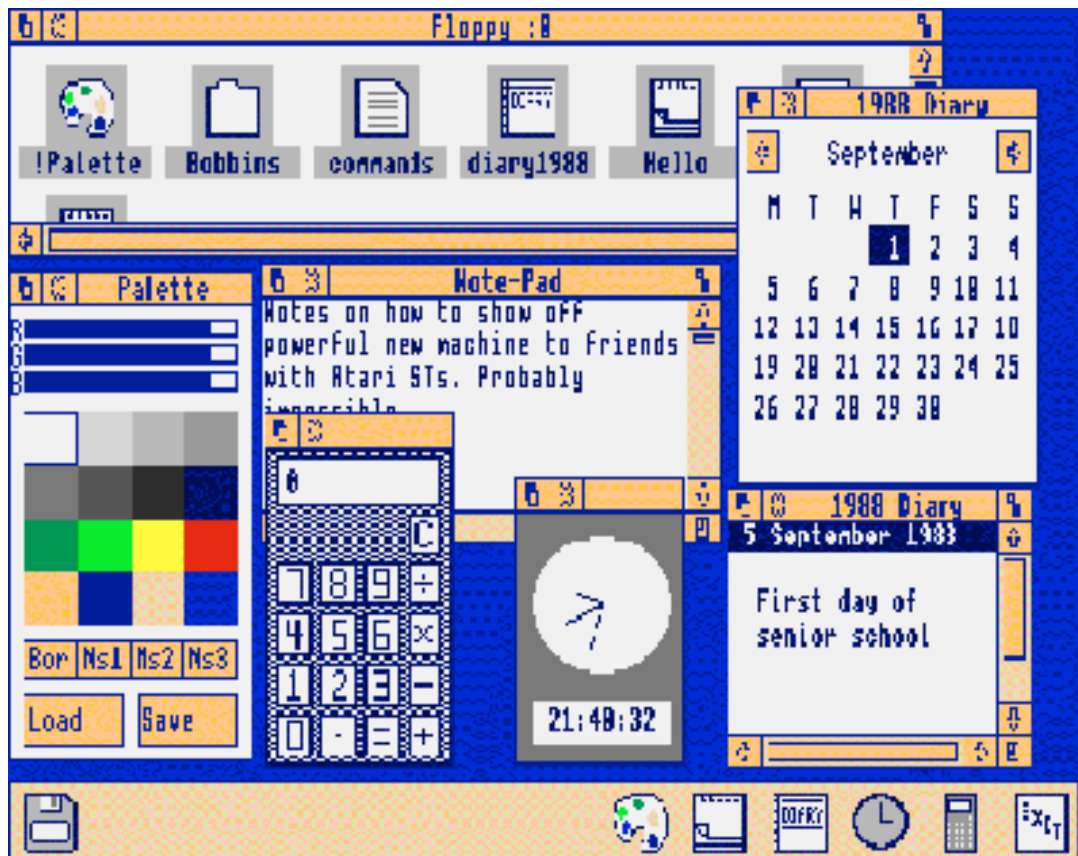
GEOS running on the Commodore 64.

Windows 2.0

Windows itself was updated to version 2.0 in late 1987, abandoning the tiled window approach in favor of the now-traditional overlapping method. This release, along with enhancement software sold by HP called NewWave, caused Apple to sue Microsoft over the “look and feel” of the GUI. Apple eventually lost this case, although HP had already withdrawn NewWave from the market by this time.

Acorn

Also in 1987, the UK-based company Acorn Computers introduced their first GUI along with what was the world’s first 32-bit RISC-based microcomputer, the Acorn A305/A310. This GUI used proportionally-sized scroll bars and introduced a new concept: a “Dock” or shelf at the bottom of the screen where shortcuts to launch common programs and tools could be kept. The GUI, called “Arthur”, also was the first to feature anti-aliased display of on-screen fonts, even in 16-color mode!

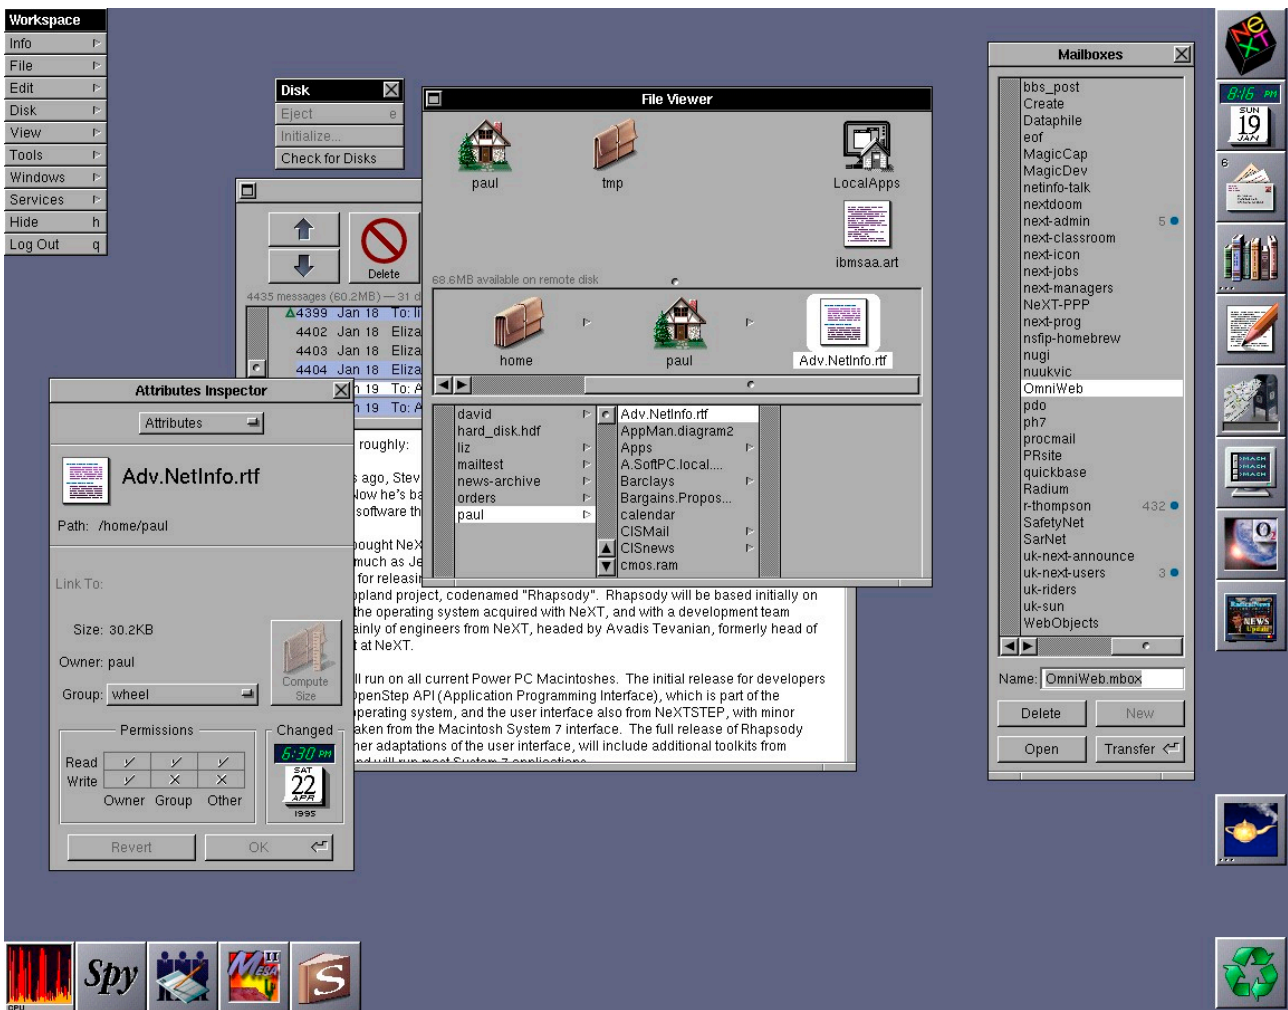


Arthur, the GUI for Acorn Computers, later renamed RISCOS.

It is important to note that many of the GUIs released in the mid-80s appear to be using fixed-width fonts, such as GEM, Windows 1.0, Amiga Workbench and Arthur. In fact all of these GUIs supported proportionally-spaced fonts in applications, but they used a fixed-width font for the system (menus and icon labels) for the sake of clarity. The reason for this is because of the low resolution of the display screens (often 640 by 200) in the hardware at the time. Even the Macintosh, with its slightly higher vertical resolution of 384 pixels, had a specially-designed system font (Chicago) that was easy to read, even if it was “grayed out” as unavailable menu options were. As screen resolutions increased, GUIs generally moved towards proportionally-spaced fonts even for the standard system text.

NeXTSTEP

1988 saw the release of NeXTSTEP, the new GUI and operating system for Steve Jobs’ NeXT computer, his first major project after leaving Apple in 1985. NeXTSTEP introduced a sharp, 3D beveled look to all of its GUI components, was the first to use the “X” symbol to indicate a close window widget, and introduced the idea of a vertical menu strip in the upper left-hand corner, which could also be “torn off” at any point so that the user could leave specific menus at any point on the screen. NeXTSTEP also had a Dock that lived on any side of the screen (but the default was on the right hand side).

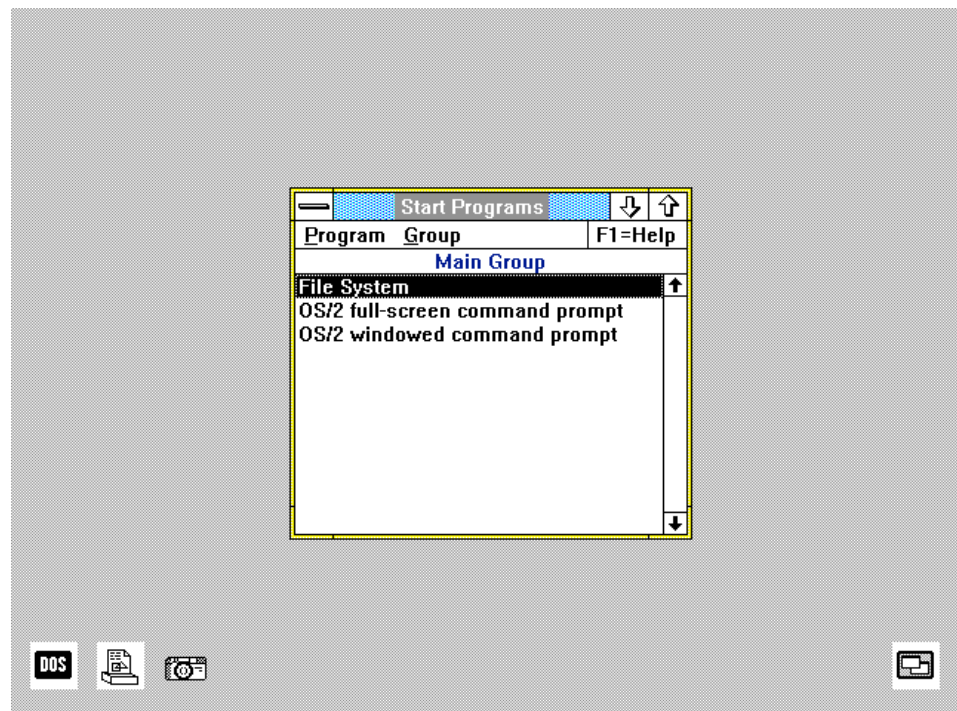


A late version of NeXTSTEP. Because it ran on NeXT's expensive custom hardware, it could demand a much higher resolution screen mode by default.*

OS/2

Also in 1988 came the first graphical version of OS/2, a project designed to replace DOS that was (at the time) a collaborative effort between IBM and Microsoft. OS/2 1.0 had been text-only, but version 1.1 came with a graphical user interface known as the Presentation Manager. This was visually quite similar to Windows 2.0.

* Der ursprüngliche Screenshot wurde ausgetauscht, da die Bildqualität zu schlecht war (red.)



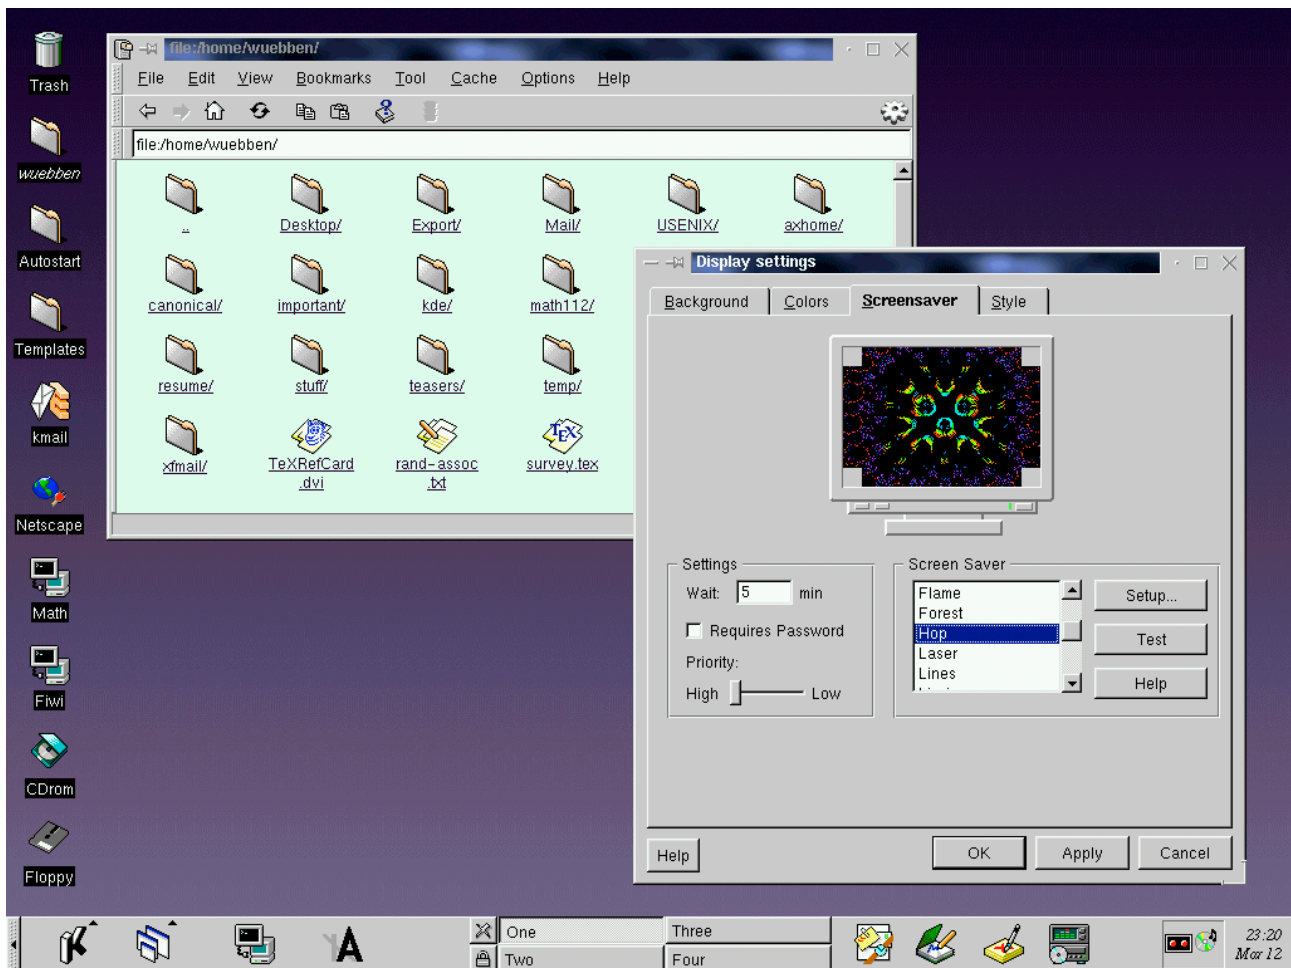
OS/2 1.1's opening screen. Note the monochrome icons — OS/2 would not support color icons until version 1.2.

X marks the spot

Just before the end of the 1980s, new GUIs started appearing on Unix workstations manufactured by AT&T and Sun (Open Look) and DEC and HP (Motif Open Software Foundation or OSF). These ran on top of a networked windowing architecture known as X, which would later be the foundation for GUIs on Linux. These were simple GUIs that attempted to mimic the appearance of Microsoft Windows but still allow access to the power of the Unix shell underneath. X also introduced a new GUI idea where merely moving the mouse cursor over a window would automatically activate it and allow the user to start typing in it.

The initial design goal of the X Window System (which was invented at MIT in 1984) was merely to provide the framework for displaying multiple command shells and a clock on a single large workstation monitor. The philosophy of X was to “separate policy and mechanism” which meant that it would handle basic graphical and windowing requests, but left the overall look of the interface up to the individual program.

To provide a consistent interface, a second layer of code, called a “window manager” was required on top of the X Window server. The window manager handled the creation and manipulation of windows and window widgets, but was not a complete graphical user interface. Another layer was created on top of that, called a “desktop environment” or DE, and varied depending on the Unix vendor, so that Sun’s interface would look different from SGI’s. With the rise of free Unix clones such as Linux and FreeBSD in the early 90s, there came a demand for a free, open-sourced desktop environment. Two of the more prominent projects that satisfied this need were the KDE and GNOME efforts, which were started in 1996 and 1997 respectively.



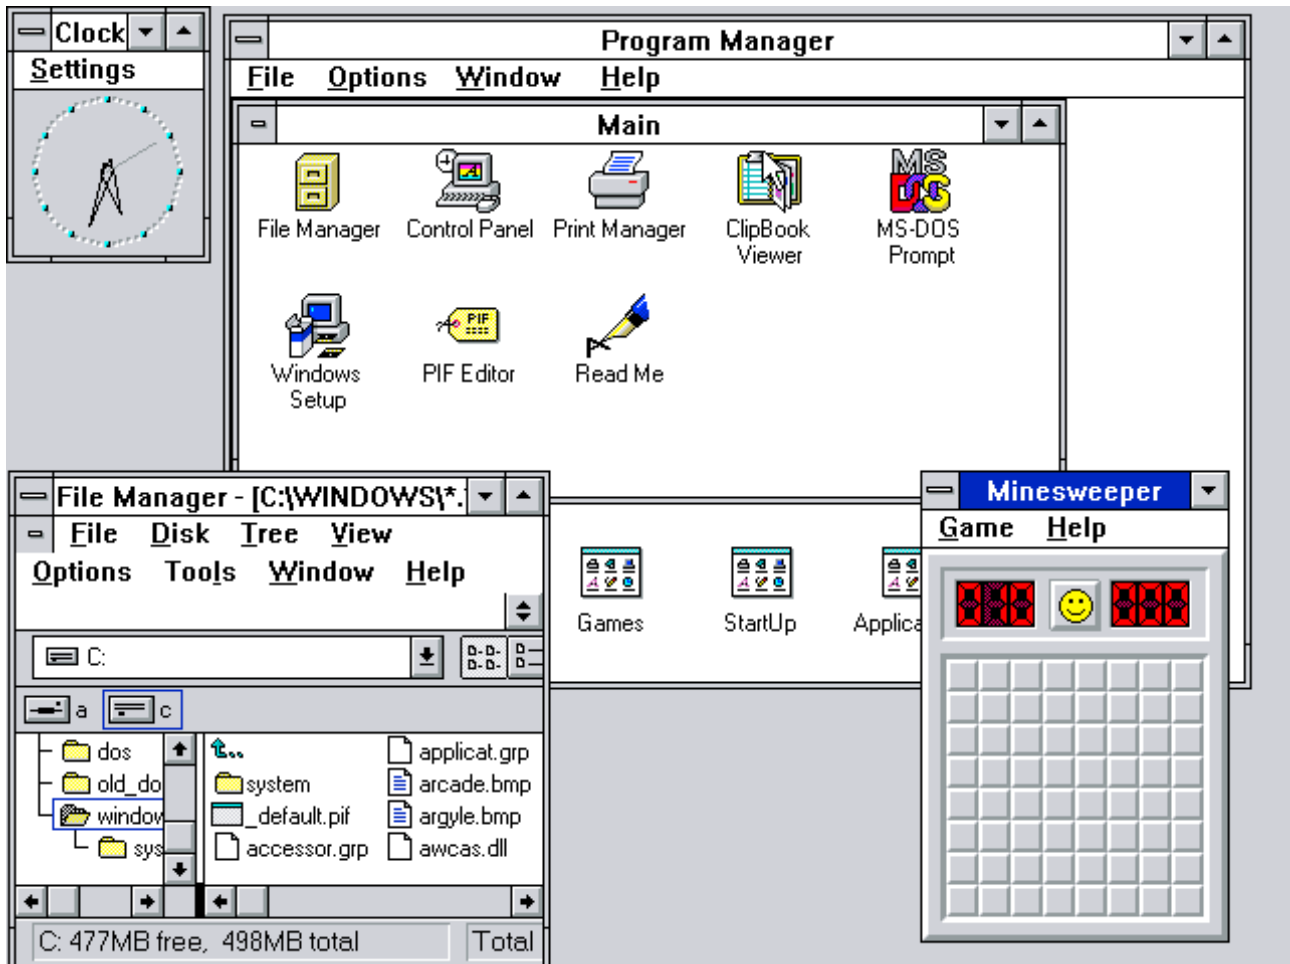
An early version of the K Desktop Environment, or KDE, for Linux and FreeBSD.

The 1990s and beyond

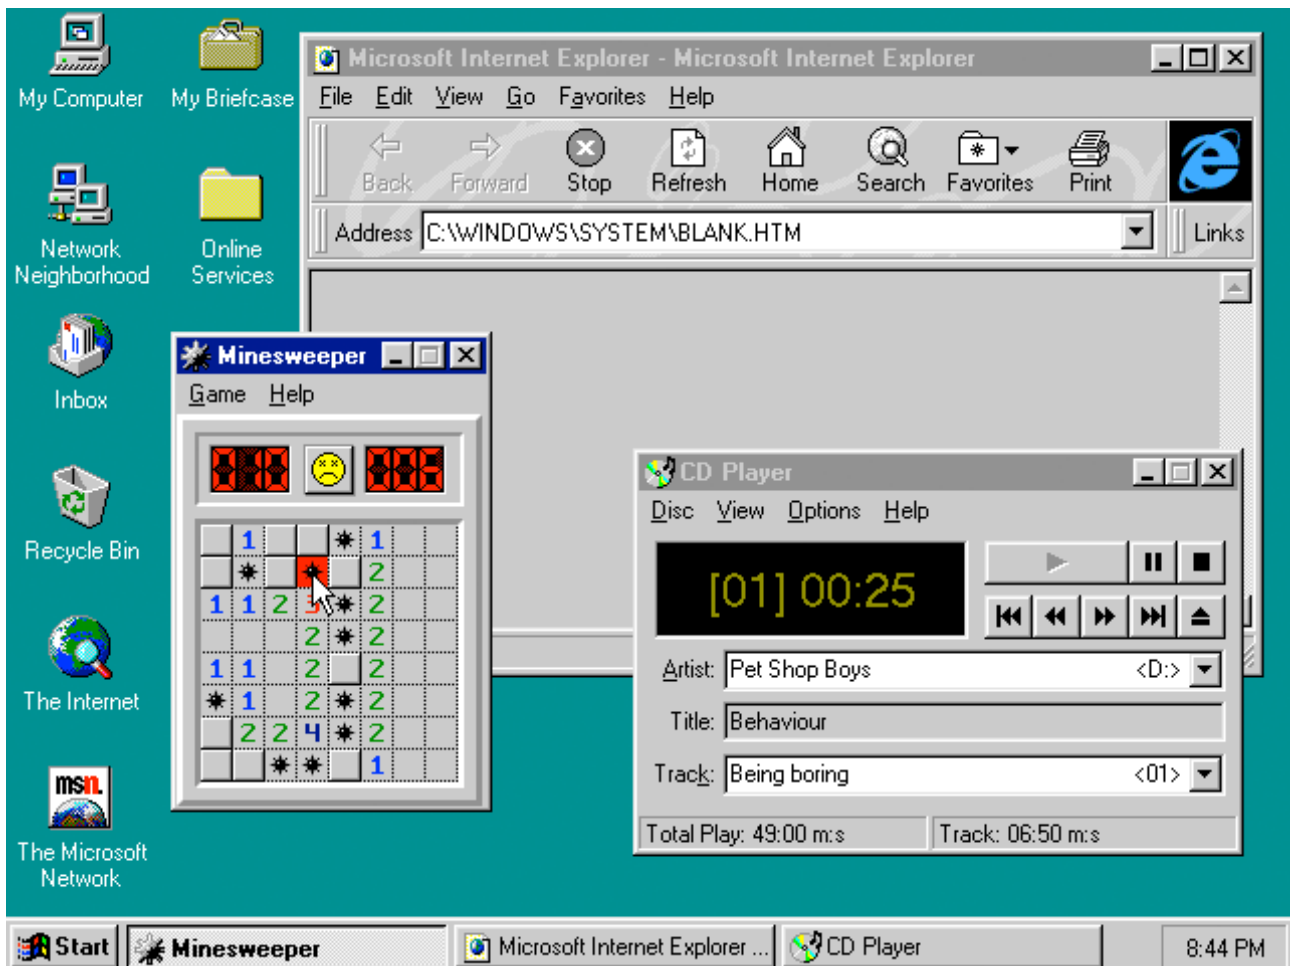
As the 90s began, other personal computing platforms fell off sharply in popularity, leaving only Windows and the Macintosh as the survivors of the GUI wars. Digital Research stopped selling GEM in 1989 and the company was bought by Novell a few years later. Atari stopped selling its ST line in 1993, the same year NeXT stopped selling its own hardware to focus entirely on the OS — they would be bought out by (or take over, depending on how you look at it) Apple in 1997. Commodore went bankrupt in 1994. Arthur, part of RISC OS on the Acorn computer, would continue to be sold until 1998.

Windows

Windows reached an unprecedented level of popularity with the release of version 3.0 in 1990 and 3.1 in 1992. While still lacking many of the features of the Macintosh (such as an icon-based file manager) it was sharp and had good looking icons, and sold millions of copies. The release of Windows 95 cemented Microsoft's lead in GUI sales, and became one of the most popular programs of all time.



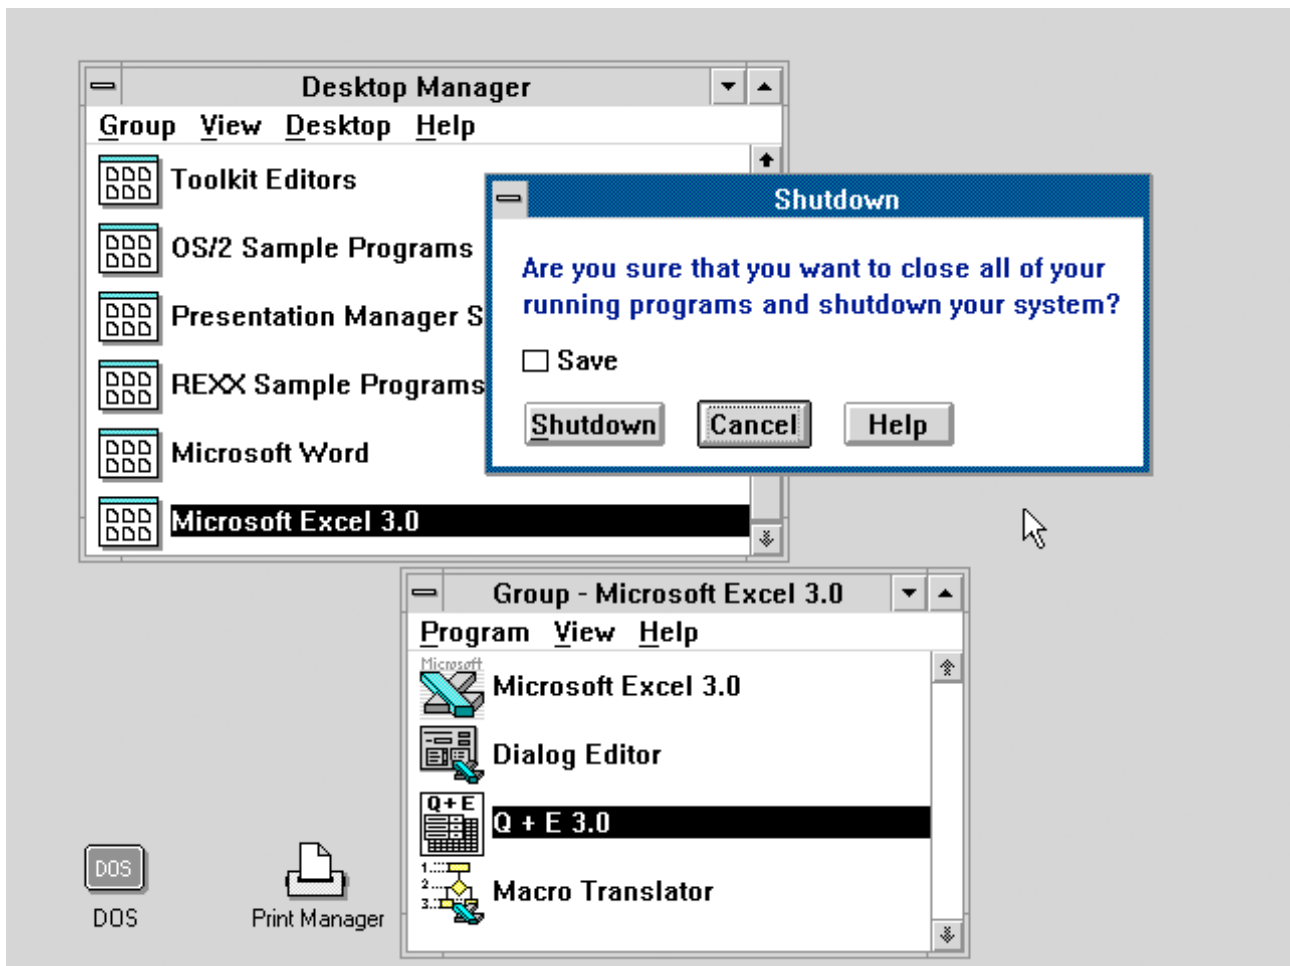
Windows 3 was the first widely popular GUI on the PC.



Windows 95 broke all sales records and became the most popular GUI of all time.

OS/2 in the 1990s

The popularity of Windows 3 and lackluster sales of OS/2 caused a much-publicized split between IBM and Microsoft over the future of both those products. IBM took over all future development of OS/2 and left Windows to Microsoft. In 1992, IBM released a new version of OS/2, 2.0, with a brand-new user interface called the Workplace Shell. The Workplace Shell took some ideas from NeXTSTEP (which IBM had licensed from NeXT but never used) and featured the ability to change the appearance of the interface by dragging and drop icons, fonts, and background colors almost anywhere on the desktop. OS/2 2.0 also had the ability, thanks to the license agreement with Microsoft, of running a built-in copy of Windows 3.1 in a virtual machine. This allowed the user to run Windows 3.1 applications either full-screen or on a window on the OS/2 desktop, although the visual differences between each interface made the latter mode somewhat disconcerting (for example, the mouse pointer would change color and shape depending on whether it was hovering over an OS/2 or Windows window).



OS/2 2.0 introduced the Workplace Shell user interface.

BeOS

Despite this apparent homogenization of software interfaces, there was still room for innovation. Windows 95 introduced the concept of the Start Menu, from which all programs could be launched, and the Task Bar where all running programs could be switched between. BeOS, introduced as part of the BeBox computer in 1995 and as an operating system for the PC in 1998, introduced the idea of "Taskbar Grouping" where tasks were sorted by application type, so each document loaded in a word processor could be found in a submenu under that word processor. BeOS also added to the idea of the Smalltalk-like title bar by allowing the user to move it left or right along the top of the window, in order for background applications to still be visible.



BeOS introduced new concepts in GUI management.

Mac OS X and Aqua

Apple did not stand still either, developing the new GUI called Aqua for their new operating system Mac OS X, itself a result of the merger with NeXT and at its core a new version of NeXTSTEP. Aqua introduced the idea of a GUI where every window was double-buffered in memory, so that any redraws happen off-screen and aren't visible (try the same "moving a Finder window over a Microsoft Word document" trick in OS X and no matter how fast your eyes are, you won't see a redraw).

Aqua also introduced several eye-candy features like minimized windows stretching and squeezing into the dock, and the concept of "sheets" where a dialog box appears to zoom right out of its attached application. In Apple's latest version of OS X, a new feature called Exposé put a new twist on application switching by growing and shrinking every open application's window in order to fit them all on one screen.



The Mac OS got a complete face lift with the release of OS X 10.0.*

Conclusions

The history of the development of the graphical user interface is a long and complicated tale. While it is easy to find individuals like Douglas Engelbart and Alan Kay who made great contributions to advancing the state of the art, the truth of the story is that the GUI was developed by many different people over a long period of time. Saying that “Apple invented the GUI” or “Apple ripped off the idea from PARC” is overly simplistic, but saying that “Xerox invented the GUI” is equally so. In fact each team borrowed liberally from all GUIs that had been created in the past, added their own unique contributions, and paved the way for other teams to move forward in the future.

* Der ursprüngliche Screenshot wurde wegen schlechter Qualität gegen einen aktuellen OS X 10.4 ausgetauscht (Quelle: Apple Computer Inc.)

Many people consider the GUI to be stagnant, differing little in its basic desktop, windows, mouse, icons, and pointer concept from the original Lisa of 1983. In some respects this is because people became familiar with the Lisa/Macintosh style of graphical interface and future projects leveraged that familiarity. However, given the extremely long gestation of the GUI from its humblest beginnings, and given that personal computer sales rose exponentially only in the mid-1990s, it is probably more accurate to think of the GUI as a slow evolution towards an ideal interface. While some attempts have been made (such as Sun's Looking Glass demo and Microsoft's 3D user interface research project) to radically change the way we interact with the GUI, the chances of these types of changes making their way to mainstream GUIs seem remote.

However, as we look forward to Longhorn and future versions of Mac OS X, we can see that although much of the core functionality of the GUI remains unchanged since its earliest debut, the potential for adding new features and modes of interaction remains limitless.

GUI development timeline

