STARS: RISE AND FALL OF MINICOMPUTERS

During the 1960s a new class of low-cost computers evolved, which, combined with a growing number of individuals programming and using them, led to wide-ranging innovations that helped to stimulate decades of growth of the entire computer industry.

Editor's note: *This month we continue a series of reprints from the IEEE Global History Network's STARS articles.[1] STARS is an online compendium of invited, peer-reviewed articles on the history of major developments in electrical and computer science and technology. Some light editing has been done to make the article more suitable for a journal publication.*

I. ESSAY

THERE is wide agreement among scholars and participants in the computer industry that a class of computers called minicomputers existed, and that it has had a significant, positive, impact on the development of computers and the industry. Reaching agreement on the exact boundaries of that class has not, however, been easy. Minimal, or small, general purpose computers were first introduced in the late 1950s, although they were not given the name "minicomputer" until 1967. By 1970, almost 100 companies had formed to manufacture these computers for new applications using integrated circuits. Minicomputers were distinguished from the larger "mainframes" by price, function, size, use, and marketing methods. By 1968, they formed a new class of computers, representing the smallest general purpose computers. While mainframes cost up to $1 000 000, minicomputers cost well under $100 000. The computer that defined the minicomputer, the PDP-8, cost only $18 000.

[1]Please refer to the STARS website for additional information and to view the full article: http://www.ieeeghn.org/wiki/index.php/STARS:Rise_and_Fall_of_Minicomputers.

Mainframes required specialized rooms and technicians for operation, thus separating the user from the computer, whereas minis were designed for direct, personal interaction with the programmer. Mainframes operated in isolation; minis could communicate with other systems in real time. In contrast with much larger mainframe memories needed for scientific calculations and business records, the first minis stored only 4096 words of 12 or 16 bits. Unlike larger computers featuring expensive input/output (I/O) devices, early minis used only a Teletype or a Flexowriter and a paper-tape punch/reader. Minis were designed for process control and data transmission and switching, whereas mainframes emphasized data storage, processing, and calculating. Minicomputers that did the most to define the new class of computers were sold to original equipment manufacturers (OEMs) that incorporated minis into larger control systems, often for industrial processes.

II. ESTABLISHING THE MINICOMPUTER INDUSTRY (1956–1964)

During the first generation of computers, those built largely prior to 1960, companies made and sold or leased room-sized mainframe computers based on vacuum-tube circuitry, mainly for record-keeping purposes. Between 1956 and 1964, groups at universities, laboratories, and entrepreneurial companies began to exploit the advantages of transistor circuitry and magnetic core memory to meet demands for computer processing by smaller organizations. They designed, assembled, programmed, and sold smaller computers that cost an order of magnitude less than traditional systems. Two small computers costing under $50 000, the Bendix G-15 and Librascope's LGP-30, were introduced in 1956 for personal, general-purpose, computation. (See Fig. 1.) Programs were prepared offline and paper tape and typewriters comprised the I/O devices. About 460 LGP-30s and 300 G-15s were leased or sold before transistorized models were announced in 1961–1962; they were the second and third most popular computers of the 1950s, behind the IBM 650.

In the early 1960s, a business plan to build and sell computers was deemed risky, expensive, and infeasible for small or new companies. IBM, and its competitors building large computers, had mastered mainframe circuit design; IBM had mastered high-volume manufacturing techniques. Most visibly, IBM dominated the market for computers. It and seven other companies were responsible for the vast majority of the computers installed worldwide. Beyond its mainframes, IBM's sales force also sold far cheaper small computers, such as the desk-sized, variable-word-length 1620

Timeline	
1956	Bendix Aviation G-15 and Librascope LGP-30 desk-size computers are marketed
1960	CDC 160 minicomputer and DEC PDP-1 are introduced using transistorized circuitry
1961	MIT develops Compatible Time Sharing System for an IBM 7094 and a DEC PDP-1
1962	MIT Lincoln Laboratory introduces LINC as interactive personal computer
1964	Computer Controls Corporation introduces DDP-116 as first 16-bit minicomputer
1964	IBM announces System/360, establishing 8-bit byte as mainframe standard
1965	IBM introduces 16-bit 1130, the first computer leased for under $1,000 per month
1965	Gordon Moore predicts increasing semiconductor chip densities, later known as Moore's Law
1967	First use of "minicomputer," following miniskirts and Austin Mini automobile
1969	Data General's NOVA lowers 16-bit processor cost with single printed-circuit board
1970	DEC announces PDP-11/20; last models are introduced in 1990
1974	Interdata sells its 7/32 as first 32-bit minicomputer for high performance applications
1979	Motorola announces 32-bit 68000 microprocessor for UNIX operating system
1981	IBM introduces its first Personal Computer using Intel X86 and Microsoft MS/DOS
1984	92 U.S. minicomputer companies in business; only seven remain independent in 1994
1985	DEC withdraws PDP-8 from the market
1985	Shared-memory multiple microprocessors are introduced by many companies
1995	Beowulf-Linux cluster operating system released for PCs and networks

for $85 000, beginning in 1959. Banks and investors assumed that IBM's industrial, technical, and commercial dominance meant that no openings existed for entrepreneurs wanting to produce computers for smaller organizations. They overlooked, however, the expanding market that the declining cost and growing capability of solid-state circuits made available. (See Fig. 2.)

Innovative transistor circuit design and manufacturing technologies were the keys to having a "first mover" commercial advantage with customers who wanted a computer's capabilities without its mainframe costs. The number of new customers was an order of magnitude greater than the 9300 computers operating in the United States in 1962.

In 1960, Control Data Corporation (CDC), Packard Bell, and Digital Equipment Corporation (DEC) all introduced small or mini computers using discrete transistors for offline processing, general computation, and process control—e.g., in power plants or with scientific instruments. CDC was a spinoff from Sperry Rand Corporation whose major asset was computer designer Seymour Cray; its 160A cost $60 000 and fit into, and on, the equivalent of an office desk. Packard Bell designed its $49 500 digital PB250 for low-cost, high-speed, hybrid computing in tandem with desktop analog computers, which was of particular interest in laboratory environments. DEC had spun out of MIT's Lincoln Laboratory in 1957, where DEC co-founder Kenneth Olsen led the circuit design for the transistorized TX-0 computer in 1956. The new company began by designing and selling transistorized circuit modules for building digital systems; it subsequently applied these techniques and products to its own computers, beginning with the PDP-1 in 1960.

Fig. 1. *A teacher holds a segment of paper tape while a student sits at the Flexowriter with his right hand on the carriage. The Librascope LGP-30 control buttons are on his right. Courtesy of the Computer History Museum.*
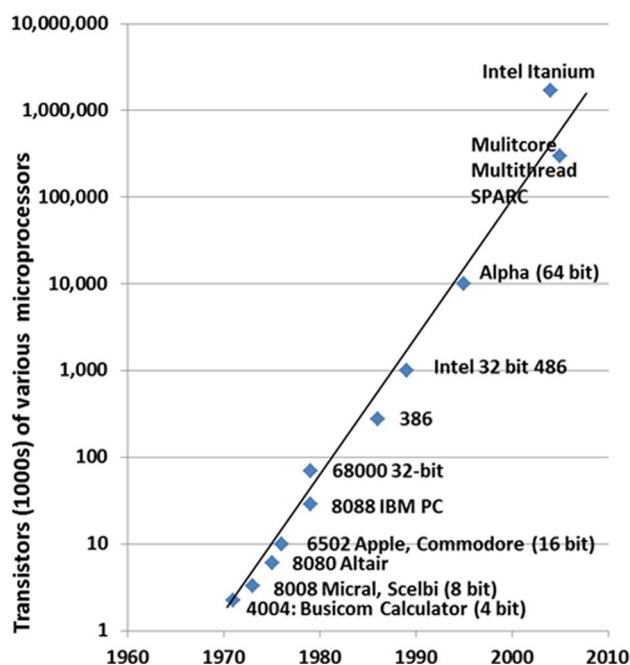
Fig. 2. *Moore's law, whereby the number of transistors on the most cost-effective chip doubles every 24 months, has enabled the creation of new computer classes: calculators, home computers, personal computers, workstations, super minicomputers, multiple processors, and multiple computers. Courtesy of the author.*

With its 18-bit, 4096-word memory, DEC's PDP-1 was priced from $85 000 to $100 000 for a standard installation. Its model name, for "Programmed Data Processor," helped especially government buyers circumvent long procurement processes intended for mainframe computers. Designed to run on standard power in a typical office, the PDP-1 contrasted with mainframes in its installation requirements, small size (17 ft²), and user-friendly interface. It was designed to connect to and support a variety of I/O devices for high-data-rate, real-time uses. For example, it had a multilevel interrupt capability to respond to real-time devices; nearly half of the 53 PDP-1s sold went to ITT for "message switching" interfaces with teletype lines. Drawing on the TX-0, PDP-1 introduced direct memory access (DMA) ports so that data from external devices could be fed directly into memory with minimal demands on the central processing unit (CPU). It had an optional point-plotting 13-in display and light pen for user interaction that was also inspired by the TX-0.

At MIT, PDP-1 contributions to software included a music compiler, program debuggers, and text editors that anticipated word processing and digital typesetting. Several MIT students wrote *Spacewar!*, one of the first interactive computer games, on a PDP-1. They became known as the first hackers.

Interactive personal uses constituted one path to timeshared computing. Concurrently with the MIT's Compatible Time Sharing System (CTSS) for the IBM 7094, DEC and Scientific Data Systems used small computers in the early 1960s—the DEC PDP-1 and SDS 930, respectively—because of their cost and flexibility in their first time-sharing systems. Between 10 and 30 users could use a computer simultaneously by sharing a large centralized computer on a millisecond-by-millisecond time-slice basis.

Wesley Clark at the MIT's Lincoln Laboratory designed another approach to personal computing in 1961–1962: the Linc, later known as the Laboratory INstrument Computer (LINC). This minicomputer was also self-contained for a researcher's solo use and, as important, inexpensive. It

had an appropriate interface, keyboard, and display, and it provided programing and data filing systems that enabled program creation, testing, and operation. Because its design was in the public domain, various organizations and individuals built 50–60 of them; DEC made and sold 21 for $43 600 apiece. (See Fig. 3.)

A major innovation was the "LINCtape" that DEC adopted as DECtape. Unlike magnetic data processing tapes, it was block addressable, which enabled a personal filing system of directories and files; it was also highly reliable because of its dual redundancy formatting and Mylar-coated magnetic layer. Until the floppy disk arrived in the mid-1970s, DECtape gave DEC a significant edge over competitors using paper tape or cards.

In 1963, DEC built the PDP-5 to interface with a nuclear reactor at Canada's Chalk River Laboratories. The design came from a requirement to connect sensors and control registers to stabilize the reactor, whose main control computer was a PDP-1. The PDP-5 was soon sold as an alternative to a custom-designed, fixed-function, hardwired control system that demonstrated the power and versatility of computing. With a memory of 4096 (4 K) 12-bit words, it cost $27 000 including I/O devices. Its 19-in cabinet contained a hand-wired backplane filled with 150 printed circuit board modules holding over 900 transistors. Upgradable and minimal in design, the PDP-5 was specifically designed to connect with external devices using a bus for both program control and direct memory access.

III. TRIUMPH OF MINICOMPUTERS (1965–1974)

By 1965, when we mark the beginning of the commercial minicomputer era, about a dozen startups and existing companies were building computers for process control (e.g., Bailey Meter, Foxboro), scientific instrumentation (Beckman Instruments, Varian), and military systems (Hughes, Raytheon).

Fig. 3. *Maury L. Pepper of the Washington University's Computer Systems Lab in St. Louis, MO, USA, developed a floating-point arithmetic subroutine on its LINC in the mid-1960s. Courtesy Maury L. Pepper.*

There was also IBM, and revenues from its 1620/1710 general purpose computers dwarfed the total revenues of the handful of companies making small computers. IBM's decision to standardize on the eight-bit byte in 1964 with the hugely successful System/360 established the standard for subsequent computers—with the implication that future word lengths would be binary multiples of eight. Its ''desk-sized'' 1130, advertised in March 1965, used IBM's compact Solid Logic Technology circuitry and 16-bit words for businesses in ''publishing, construction, finance, manufacturing and distribution.'' A minicomputer in all but given name, up to 10 000 1130s leased for $695 per month or sold for $32 280 each, supported by programs for applications in the above industries. Its 4096-word core memory could be connected to IBM's magnetic disk storage unit; users could input and output data and programs through paper tape, punch cards, or a printer. (See Fig. 4.)

A month later DEC advertised the 12-bit-word PDP-8, widely acknowledged as the definitive mini but not described as such until 1967. Sold for only $18 000, it was the least expensive computer on the market, undercutting IBM and other makers of small computers. At its introduction,

the PDP-8 had a cost and first-mover advantage based on proprietary design, diode-transistor logic (DTL), and printed circuit boards. These were interconnected using automatically wire-wrapped panels to lower costs. Truly minimal, it occupied only half of a standard electronics equipment cabinet; yet, over its 20-year lifespan, users applied it for process control, instrumentation, message switching, typesetting, business computing, personal computing, and, eventually, word processing. (See Fig. 5.)

The "8" was also used for time-sharing. At Carnegie-Mellon University in 1968, DEC's TSS/8 project demonstrated a low-cost, interactive system for educational use, data entry, and terminal multiplexing.

Most importantly, DEC's willingness to share design details encouraged customers to develop their own software and applications. It pioneered sales to original equipment manufacturers (OEMs) that integrated minicomputers into their proprietary electronic systems. The majority of over 50 000 PDP-8s, the bestselling computer of its era, were bought, adapted, and resold by OEMs. In addition, the minicomputer pioneered

Fig. 4. *Two secretaries work at the keyboard and printer of IBM's 1130 minicomputer in this promotional photo, c. 1965. Courtesy IBM Archives.*

Fig. 5. *DEC promotional photo of woman operating a PDP-8 in an office environment, c. 1965. Courtesy of the Computer History Museum.*

and created departmental computing within larger organizations, changing the culture of institutional computing. Now smaller groups could buy, install, and maintain their own computers for particular uses. In this way, computing began to migrate from operation in a single central facility to use based on functional needs defined by a given department in a government, a corporation, or a university.

In September 1965, Computer Control Company introduced its 16-bit word mini to conform to IBM's 8-bit standard; henceforth most computers would be implemented with modular 8-bit word lengths, i.e., 8, 16, 24, 32, and 64 bits. The $28 500, one-cabinet DDP-116 designed by Gardner Hendrie began competing with the PDP-5 and IBM's 1130. A year later, however, Honeywell purchased the company and introduced its version of the DDP-516, marketed as the H-316 in 1969. Bolt, Beranek, and Newman (BBN) adopted a militarized version that year as the ARPAnet's Interface Message Processor for packet switching and timesharing with user terminals and computers. In the commercial sphere, however, Honeywell did not structure

itself well for the emerging mini market.

The structure was almost as awkward at Hewlett-Packard (HP). David Packard believed that computers would be invaluable in processing data from the electronic instruments that HP sold. His business partner William Hewlett disagreed, leading to several false starts for minicomputer innovation in their company. Nonetheless, after abandoning an attempt to buy DEC in the early 1960s, HP created a new division whose staff drew on the rights to Data Systems Inc.'s DSI 1000 computer, HP's new integrated circuits, and a concept for an instrument controller. The result was the 16-bit, $22 000 2116A, introduced in 1966 to interface in real time with a wide range of instruments. Follow-on models were bought more for timesharing than as controllers and sales persisted into the 1980s. Despite this unexpected bump in sales, HP canceled its pioneering 32-bit "Omega" prototype in 1969. Instead engineers reworked it as a 16-bit general purpose mini with a "multiprogramming executive" (MPE) operating system more versatile than the 2100's. This "Alpha" project was

marketed as the 3000, a successfully evolving series which, after a disastrous introduction in 1972–1973, HP sold and serviced into the 2000s.

By 1970, about half of the new minicomputer companies had established themselves using large-scale integration (LSI) and bipolar integrated circuits to produce 16-bit and 8-bit computers. Data General's 1969 16-bit Nova was noteworthy because its clever design put the processor on a single, large, printed circuit board, eliminating the backplane wiring and reducing costs. Led by former DEC manager Edson de Castro, Nova's reliable architecture and automated packaging made DG a major minicomputer competitor.

One result was that it accelerated DEC's introduction of its 16-bit PDP-11. (See Fig. 6.) This minicomputer was the industry benchmark until the early 1980s; the final single-chip models were introduced in 1990, with about 200 000 sold altogether. Its popularity was based on ease of programming, thanks to the use of general registers and an orthogonal instruction set, and a flexible I/O structure that enabled connections to proprietary peripherals. In addition, the "11s" were managed and controlled by any of seven user-defined operating systems, for real-time process control, general purpose timesharing, interactive hospital patient record management, and single-user personal computing.

DEC's commercial dominance through the PDP-8 and -11 often meant that exploiting a niche for product differentiation was the key to its competitors' survival in the 1970s. Thus, many minicomputer firms used proprietary architectures and their own software. These typically included an operating system, FORTRAN compiler, and program development environment. At the low end Computer Automation introduced its 8-bit and 16-bit "Naked Minis" in 1972 at $1450 and $1995, respectively, for OEM buyers. General Automation started in 1968 with its 8-bit SPC-12 "Automation Computer" and in 1975 introduced the 18/30, which was compatible with IBM's

IV. MINICOMPUTERS AND MICROPROCESSORS (1970–1984)

A computer architecture's potential performance is determined predominantly by its word length—the number of bits that can be processed per access—and, more importantly, memory size, or the number of words that it can access. Most architectures can be extended in some fashion, but the speed at which chip companies innovated microprocessors shortened their lifespans. The 1970 DEC PDP-11 claimed to provide an adequately large address space of 16 bits to address up to 64 kB. By 1975, however, DEC had added a 4-bit memory extension to its PDP-11s to access megabyte memories for timesharing systems. At the same time, other companies were pushing into terrain once reserved for mainframes. Prime introduced a clone of the Honeywell H-516 in 1972 using 32-bit words. With special memory management or paging software, it had a unique advantage for running large programs. Interdata introduced its 32-bit 7/32 using microcoding in 1974 to implement a complete Instruction Set Architecture, especially floating point. This enabled the company to be a significant supplier for real-time embedded systems, e.g., in-flight simulators, CAT scanners, and power plants. The 8/32 "Megamini," introduced a year later, generated the imagery for the original film *Tron* (1982).

In 1978, with the VAX-11's 32-bit architecture in place, DEC asserted that 32-bit addresses accessing 4 GB would be adequate in a VAX-based strategy of distributed computing. This covered shared or networked computers as well as single microcomputer terminals. Over the VAX's 20-year life, approximately 140 models were introduced using transistor–transistor logic (TTL), emitter-coupled logic (ECL), and custom complementary metal–oxide–semiconductor (CMOS) single-chip technologies. Doubling the memory size to 32 bits, however, took the mini into direct competition with smaller mainframes that used 32-bit architectures,

Fig. 6. *Dennis Ritchie (standing) and Ken Thompson, inventors of the UNIX operating system at Bell labs, work at a PDP-11 in the early 1970s. Courtesy Alcatel-Lucent.*

1130 and 1800 process control computers. Microdata began implementing microprogramming to ease user customization in its minicomputers in the early 1970s. Others exploited AMD's 2900 IC family of registers, data paths, and microprogrammable control components to implement proprietary architectures from 1975.

Although the minicomputer was created for technical computing (real-time process control, data acquisition, simulation, and interactive use), commercial computing (data entry, word processing, spreadsheets) was significant. When IBM minicomputers are considered, commercial use undoubtedly dominated. Interestingly, there is no record of IBM referring to any of its computers as minicomputers. The media and other industry observers had no such reticence. For example, when IBM announced its Series/1 Model 3 and Model 5 computers for purchase at prices from $10 000 to $100 000, the *Palm Beach Times* in Florida headlined its November 17, 1976, issue, "IBM introduces new minicomputer." The company continued to produce and service "mid-

range" computers from the System/38 in 1979 to the AS/400 in 1989 and Power Systems servers of 2008.

Others joined in. Four Phase Systems introduced the System IV/70 for database access in 1970 using the company's custom MOS AL1 microprocessor. Viatron attempted to enter the low-cost, data entry market based on custom, very large-scale integration (VLSI) chips that it could not manufacture. Wang Laboratories, a company known for calculators, began evolving in the early 1970s into a major company for word processing and office applications in competition with IBM. Other companies including Burroughs, NCR, and Univac comprised lesser parts of the commercial market. The financial industry, government, and manufacturers used a significant number of minis from DEC, DG, and Prime, among others. Prime's success stemmed from founder William Poduska's training in early timesharing at MIT, which the company's staff applied to a series of 32-bit minicomputers using the PRIMOS operating system.

including the IBM System/360. It also forced smaller competitors to upgrade their products, which in the case of Data General resulted in the MV/8000 Eclipse superminicomputer in 1980, and a Pulitzer-Prize-winning book on its innovation, Tracy Kidder's *The Soul of a New Machine* (1981).

In 1971, Intel introduced its 4004 chip, the first microprocessor for sale as a component, albeit one with a 4-bit data path used to operate a calculator. Taking that chip as the initial data point, the history of computing in terms of word length and address size can be observed (see Fig. 2). Increased transistor density enabled Intel's 8-bit 8080 for the MITS Altair 8800 (1974) and Zilog's Z-80 for Radio Shack's TRS-80 (1976); and the MOS 6502 16-bit microprocessor for home computers from Apple (Apple II, 1977) and Commodore (PET, 1977). By 1981, the 16-bit Intel 8086 had been adopted by IBM for its PC. IBM's adoption of Intel's microprocessors and Microsoft's MS/DOS operating system in 1981 established the IBM PC or "Wintel" industry standard for the new personal computer class. Consequently straightforward programs and applications migrated immediately to Wintel PCs.

The steady improvement of microprocessor capacity contributed to the growing identity crisis for minicomputers and the companies that made them. By 1976, a PDP-8/A sold for $1317 as it was reconstructed with an Intersil 6100 CMOS microprocessor. The Motorola 68000 microprocessor, introduced in 1979, struck directly at the proprietary minicomputer market with its 32-bit word length, large address space, and software components that could be easily assembled to form entire computer systems. It powered Unix-based workstations and servers before Apple incorporated it in the Macintosh PC (1984). (See Fig. 7.)

The 32-bit-based microprocessors, however, also enabled startup companies to compete with established minicomputer firms along three evolutionary paths. (See Fig. 8.) These

Fig. 7. *Motorola's 68000 series microprocessor is less than 2.5 in or 62 mm long, but the 16/32-bit hybrid architecture of the computer chip under the square heat spreader makes it the equal of classic minicomputers. It is still in production in 2012. Copyright Mark Richards. Courtesy of the Computer History Museum.*

were, first, cheaper microcomputers assembled from board-level components; second, workstations—e.g., Apollo, SGI, Sun Microsystems—to address both a new market and compete with minis; and third, scalable, multiple-processor computers. The latter provided both higher performance and wider availability to users than custom-designed minis with proprietary architectures.

V. DECLINE OF THE CLASSIC MINICOMPUTER (1985–1995)

One result of this diversification was that the minicomputer industry also evolved: from one vertically integrat-

ed with proprietary architectures, TTL implementations, and proprietary standards to a horizontally disintegrated industry with standard chips, boards, peripherals, and operating systems. The price segment that helped define minicomputers, $10 000 to $100 000, also expanded by an order of magnitude.

While the demise of classic minicomputers was clear by 1985, companies continued offering them until the early 1990s, when the firms went bankrupt or were acquired by more astute competitors. (See Table 1.) Wang declared bankruptcy in 1992. Compaq bought DEC in 1998, and HP acquired Compaq in 2002. EMC turned Data General into a data storage business in 1999.

Beyond the microprocessor, other technological changes drove this industrial transformation by lowering the barriers to entry in computer design and production. Low-cost modular packaging allowed systems to be built with less effort in design and more focus on production-line assembly. Developed around 1986, VMEbus modules of processor, memory, and I/O were simply plugged into a standard back panel. In addition, the unlicensed availability of AT&T's UNIX operating
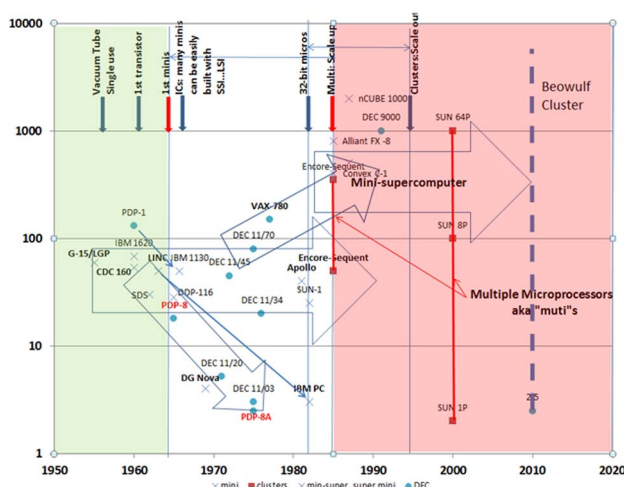
Fig. 8. *The minicomputers created in 1965 with integrated circuits led to three evolutionary paths characterizing a computer class's trajectories plotted logarithmically by price (in $1000s not adjusted for inflation) over time: declining price, which allows the formation of a new computer class (Bell's Law); continued development of integrated circuitry at constant price that allows more performance (Moore's law); and performance and functional improvements that address new uses and potentially conflict with higher priced systems.*

Table 1. 92 U.S. minicomputer companies, 1968-1985
(created by the author with help from colleagues over many years)

The following list includes all firms making general and special-purpose minicomputers for real-time processing, communications, business etc., sold to end users through OEMs, and bundled for process control and testing. It does not include scores of military, AT&T, European, and Japanese computers and processing systems developed for niche markets.

49 started up and retained autonomy or died
    2 grew at significant rates and continued to grow: Data General (1969), Prime (1972)
    8 grew at diminished or declining rates, or found small niches
    39 ceased to manufacture
10 merged with larger companies
8 existing computer companies built minicomputers
    2 grew rapidly: Digital Equipment Corporation (1960), IBM (1965)
25 existing companies built minicomputers for their own use
1 formed a division, Dymec, around an acquisition and grew rapidly: HP (1966)

system (OS) eliminated the most significant software barrier by providing an essentially free OS from the company or the University of California, Berkeley. Thus, anyone with a bench on which to plug components together and load software was a potential computer company. Consequently, classic minicomputer manufacturers and their new rivals evolved commercially to adapt to the new technologies in at least one of three ways.

First, the root cause of the classic mini's demise was that CMOS microprocessor design progressed more rapidly than the bipolar semiconductor architecture traditionally used in mainframes and minis. Most companies misjudged the transition. In 1991, MIPS Computer Systems' 64-bit R4000 microprocessor brought what had been supercomputer processing power in 1975 to the heirs to minicomputers: desktop workstations. These were the functional successors to the desk-sized 12-bit minis of the late 1960s. DEC's investment in its CMOS-based Alpha architecture provided competitive microprocessors for workstations into the mid-1990s.

Second, 1985 marked the introduction of multiple CMOS microprocessors that plugged into a common bus for shared memory access, creating a "multi." A multicomputer covered a range of costs and performances

greater than a family of computers based on individual models. These scalable computers were significantly more powerful than minis at an equivalent price. For less than $100 000 a multi could support 50–100 users connected via desk terminals, rivaling the cost of personal computers while providing more networked capability. At least six companies exploited the multiprocessor structure. DEC's VAXcluster was the forerunner of IBM's 1990 Sysplex and the enormous clusters of machines that now comprise cloud computing.

Third, with the creation of custom CMOS, multiple microprocessor architectures, and the stimulation of DARPA's Strategic Computing Initiative, over 40 companies responded to the call for high-performance, lower cost computers for technical applications. A handful of minis priced between $100 000 and $1 000 000 to address larger problem applications and shared uses thus became "super-minicomputers" or "Crayettes," reflecting their focus on the low end of Seymour Cray's supercomputer class. None of the companies involved lasted ten years, although Alliant, Ardent, Convex, and Stardent delivered significant numbers of these computers.

Cluster and cloud computing eroded traditional minicomputer roles and applications. In 1995, the Linux-based

Beowulf Cluster software was introduced to support commodity computing nodes and network switches. The cluster advantage—aside from the ability to scale out to any size defined by budget—was that standardization took place around the software interfaces. This enabled higher level libraries and programs, like MatLab and Mathematica, and systems for computational fluid dynamics, computational chemistry, etc., to be ported and run in a relatively standard environment. With the advent of cloud computing in the early 2000s, users could simply rent an arbitrarily large array of computers on an hourly basis.

The acquisitions and bankruptcies that eliminated the minicomputer companies did not mean the end of computers that powered mini applications. Instead, by 1995, what had been provided through proprietary and unique architectures and software were delivered on scaled-up multiprocessors or scaled-out computer clusters: i.e., the cloud or user data centers, using Intel x86, IBM PowerPC or Sun SPARC computers, controlled by Windows or UNIX variant operating systems. Ironically IBM emerged as a corporate heir to the minicomputer industry. It innovated a series of midrange "Vax killer" computers, including its 9370 "super-mini computers" of 1987 and 1989's AS/400 series. (See Fig. 9.) The latter

Fig. 9. *IBM promotional photo of office staff working with an IBM AS/400 Model 10 network, c. 1989. Courtesy IBM Archives.*

had its origins in IBM's System/38 and System/36 "midrange" computers or minicomputers, and it enjoyed support into 2008 as the eServer iSeries and then the System I, high-end enterprise servers for cloud computing.

In popular and media terms, the technological heir in 2012 to the minimal computers of the past is the $35, 1.2 oz., Raspberry Pi circuit board minicomputer from the eponymous U.K. foundation. Designed to inspire young programmers, it runs on a 700-MHz processor with 256 MB of RAM, and in the space of a credit card carries six ports for Internet, USB, and audiovisual connections. In addition Chinese manufacturers are selling the $74, thumb-sized, MK802 mini- or microcomputer with a Linux-based Android operating system, 1.5-GHz processor, 512 MB of RAM, and multiple ports, as well as a WiFi antenna. The unexpected popularity of the Raspberry Pi with hobbyists suggests that a new class of computers has arrived, separate from the highly popular smartphones.

VI. MINICOMPUTER DEFINITION AND CLASS CONTRIBUTIONS

Those who seek the definition of minicomputers from Wikipedia or other sources will learn that they represent a class of smaller computers that evolved in the mid-1960s, based on the first integrated circuits, and sold for much less than mainframe and midsized computers from IBM and its direct competitors. They will also learn that soon after Intel's 4004 single-chip CPU appeared in 1971, a minicomputer came to mean a machine that lay between mainframes and microcomputers in cost and performance.

These definitions may satisfy many, but they lack precision and nuance. The first computers to be called minicomputers had other unique features. For example, they were not designed primarily for typical business and scientific tasks; instead they were designed for functions such as process control, data-path switching, and communications while handling related tasks such as error correction. Equally important, they were primarily sold as components to OEMs that packaged them into larger products sold to end users. It is this process that created new markets and helped to define the new class of computers.

Also of concern is the use of the term minicomputer—following the development of single-chip CPUs—to describe computers between "the smallest mainframe computer and the microcomputer." These newer computers were sufficiently different, and competed in a sufficiently different market, that they represented a new class of computers. It is unlikely that we can convince the new generation of historians and technical writers to give this new class of computers a name other than minicomputer; but perhaps we can convince them to refer to the earlier minicomputers as "classic minicomputers."

By distinguishing between the classic minis and newer minicomputers, we can illuminate the significance of rapid changes in technologies and market conditions. Indeed the rise and fall of the classic minicomputer industry is a good example of Bell's Law of the Birth and Death of Computer Classes. Each new class is enabled by new technologies at a lower price, thereby establishing new uses and new standards that result in new industries. Over time, the process can repeat itself to eliminate an industry. Thus, by 1985, inexpensive, powerful microprocessor components and standard software had become broadly available for personal computers, workstations, and scalable multiprocessor systems, and the classic minicomputer class had largely disappeared.

Important contributions of the classic minicomputers have been

mentioned throughout this essay. Nevertheless, it is useful to remind ourselves that their primary long-term contribution was to foster innovation throughout the computer industry and thus contribute to its rapid growth.

Sold as components to original equipment manufacturers, classic minicomputers provided the opportunity for businesses to create new solutions for process control, manufacturing, engineering design, scientific experiments, communication systems, and many more. The original minicomputers also pioneered and created departmental computing within larger organizations. Now smaller groups could buy, install, and maintain their own computers for particular uses. In this way, computing began to migrate, from operation in a single, large, central facility, to use based on functional needs defined by individual departments within a government, a corporation, or a university. This provided flexibility and innovations at all levels. Perhaps the most important contribution was the diversification of programming itself. The original minicomputers made it possible for more people to become involved in programming, thus increasing the rate of innovation in the art of software, including the improvement of user interfaces needed for personal computing. ∎

GORDON BELL

ABOUT THE AUTHOR

Gordon Bell (Member, IEEE) is a Microsoft Corporation Principal Researcher and former Digital Equipment Corporation Vice President, where he led the development of the first mini- and timesharing computers. At the NSF's Directorate for Computer & Information Science Engineering, he initiated the National Research and Education Network. He has published extensively about computer architecture, high-tech startup companies, and life-logging.

Mr. Bell is a member of the Association for Computing Machinery (ACM), the American Academy of Arts and Sciences, and the National Academies of Engineering and Science. He received the 1991 National Medal of Technology and is a founding trustee of the Computer History Museum in California.