Besting Apple in the Business Market

The PC's initial advantage over the Apple was finesse rather than raw computing power. Its Intel 8088 processor, a cut-down version of the 16-bit 8086, was theoretically more powerful than the 8-bit processors used on earlier personal computers, but *Byte* observed that in practice "the necessity of funneling all data through an 8-bit path degrades the 8088's performance to the point where it is more like a fast 8-bit microprocessor with an extended instruction set than it is a 16-bit microprocessor."[31]

IBM's biggest success with the PC came in the business market. Larger companies, whose data processing departments were skeptical of Apple, had strong relationships with IBM. An IBM PC was comparable in cost to a video terminal. IBM even offered a special version of the PC configured with special hardware and software to mimic the 3270 series video terminals it sold for use with its mainframes.[32] But many corporate PC purchases were driven by individual initiatives, not official policies. Data processing staff, wedded to Wang word processors or IBM mainframes, were losing control. In December 1982, *Time* magazine named the computer its "Machine of the Year."[33]

The IBM PC faced little competition from Apple. The Apple II was four years old and optimized for computer hobbyists. For example, IBM's PC displayed a full eighty columns of upper- and lower-case text. Rather than add comparable features to the Apple II, Apple had incorporated them into its more expensive, and incompatible, Apple III, launched in 1980. Apple's rush to get the machine to market had combined with its desire to squeeze new capabilities onto a small motherboard to make its first batch of machines spectacularly unreliable. Apple finished up recalling and replacing the entire production run.[34]

This fiasco killed the Apple III's prospects, but Apple didn't release a major upgrade to the Apple II until the IIe arrived in 1983. Steve Wozniak later wrote, with a touch of overstatement, that "about the only salary Apple spent on the Apple II during that period—1980 to 1983—was on the guy who printed the price lists."[35] Over a ten-year lifespan the IIe sold millions of copies, many times the volume of earlier models. The new design lowered manufacturing costs enough to keep Apple afloat through the mid-1980s in its core home and education markets, whereas its useful but modest enhancements did little to challenge IBM for business use.

The 8087 and IEEE Floating Point

Doing engineering calculations or financial modeling cost a lot less with a personal computer, such as the Apple II, than with a mainframe or timesharing system. But only small jobs would fit into its limited memory and run at an acceptable speed. Complex models still required big computers. That began to change with the IBM PC. Even the original IBM PC could be expanded to much larger memory capacities than the Apple.

The other big difference was floating-point support. Since the 1950s, capable floating-point hardware support had been the defining characteristic of large scientifically oriented computers. The 8088 used in the original PC did not support floating point and its performance on technical calculations was mediocre. But every PC included an empty socket waiting for a new kind of chip, the 8087 *floating-point coprocessor*. The 8087 was the first chip to implement a new approach to floating point, proposed by William Kahan and later formalized in the standard IEEE 754. Its adoption by firms including DEC and IBM was a major advance for scientific computing, for which Kahan received a Turing Award. Code, even in a standard language like Fortran, had previously produced inconsistent floating-point results when run on different computers. According to Jerome Coonen, a student of Kahan's who managed software development for the original Macintosh, this standardization on robust mechanisms was a "huge step forward" from the previous "dismal situation . . . Kahan's achievement was having floating point taken for granted for 40 years."[36]

The 8087 was announced in 1980 but trickled into the market because it pushed the limits of Intel's production processes. Writing in *Byte*, Steven S. Fried called it "a full-blown 80-bit processor that performs numeric operations up to 100 times faster . . . at the same speed as a medium-sized minicomputer, while providing more accuracy than most mainframes."[37] The 8088 itself had only 29,000 transistors, but its coprocessor needed 45,000 to implement its own registers and stack.

Assembly language code and language compilers had to be rewritten to use special floating-point instructions, executed in parallel with whatever the main processor was doing. Scientific users quickly embraced the 8087, which made the PC a credible alternative to minicomputers. Fried had promised that "the 8087 can also work wonders with business applications," but software support was limited. Even Lotus-1-2-3, which existed only to crunch numbers, did not utilize it. Fried began a business selling patches to add coprocessor support to such packages.[38] Over time, IEEE-style floating point became a core part of every processor. By the time Intel launched the 80486 in 1989, its factories were just about able to manufacture this one million transistor chip with its built-in coprocessor. Software developers, particularly video game programmers, began to use floating-point instructions. By the late 1990s, PC processors competed largely on the strength of their floating-point capabilities.

The PC XT and the PCjr

The evolution of the PC reflected its success as a higher-end, business-oriented machine. In 1983 IBM launched the PC XT, which dropped the cassette interface and other features aimed at home users. Ambitious PC users had quickly run out of expansion slots. The XT added three more, plus sockets for a full 256 KB of RAM, which