From: Charles Ferguson and Charles Morris, Computer Wars: The fall of IBM and the future of global technology (1993)

1

Coloring the World Blue

When IBM announced the System/360 series of computers in 1964, *Fortune* magazine called it a "$5,000,000,000 gamble . . . possibly the 'riskiest business judgment of modern times.' " The name "360" came from the 360 degrees in a circle, because IBM intended the 360 to take over the *entire* world of computing—business, science, defense, everything. It was the kind of daring, "bet the company" gamble that few companies had the guts, the vision, or the resources to make. But it was not the first time IBM had rolled the dice so boldly, and it was characteristic of the style that had made IBM the world's outstanding company—the most admired, and most feared, in any industry.

The modern IBM, the company that emerged at the end of World War II as the world's leading business punch-card company, was very much the creation of Thomas J. Watson, Sr., universally recognized as one of the world's greatest sales and marketing geniuses. It was Watson who focused a small scales and measuring device company on solving business accounting problems, changed its name to International Business Machines, and piloted it to Fortune 500 status, able to hold its own in its chosen arena with much more powerful and technologically grounded giants like General Electric and Remington Rand. The secret of Watson's success, a principle that was ground into the very bones of IBM's salesmen, was that IBM would sell machines by solving problems; IBM would win if its accounting machines truly helped its customers' businesses run better. IBM's understanding of its clients, its commitment to customer support, the dedication to quality, the fanatical devotion to deadlines and delivery schedules, were legendary.

IBM was regarded as a technology lightweight at the end of the war, even though it had dipped a toe into electronic computing and had collaborated with Harvard University to build the Mark I, a giant electromechanical calculator—or a "robot brain," as the press called it—to assist in wartime codebreaking. In 1947 Watson had his engineers build the Selective Sequence Electronic Calculator, a 120-foot-long Rube Goldberg monster that stood in a huge glass enclosure in the IBM lobby in New York City, available to any scientist for free. But the world's leading computer company was Rand, and the name of its big machines, UNIVAC, became a household word when Edward R. Murrow's 1952 election night newscast used a UNIVAC to forecast the winner in the Eisenhower-Stevenson presidential race, with the big computer's whirring tape drives and blinking console looming behind Murrow like an alien presence on the CBS set.

As American business boomed at the end of the war, IBM's customers began demanding faster and bigger calculating machines to keep pace with the headlong expansion of their sales and territories. IBM was already selling electronic business calculators by the end of the 1940s—desk-sized machines that used vacuum tubes to do arithmetic and could compute a payroll ten times faster than the punch-card readers could feed in the data. In 1952 the company produced a vacuum-tube-based scientific computer, the 701; two years later, facing a chorus of demand from its punch-card customers, it introduced the 702, an electronic computer specifically

designed to replace accounting department punch-card machines. As Tom Watson, Jr., took control of the company in the mid-1950s, he consciously set out to push IBM into the newest electronic technologies. He recruited Emanuel Piore, head of the Office of Naval Research, as chief scientist, and increased research spending from about 15 percent of net income in the 1940s to 35 percent in the 1950s and to 50 percent by the 1960s and 1970s. By the 1960s, IBM's computer R&D budget was bigger than the federal government's.

The seminal event in postwar electronics was the invention of the first useful solid-state semiconductor electrical device at Bell Labs in 1948. The inventors, William Shockley, Walter Brattain, and John Bardeen, shared the 1956 Nobel Prize for their achievement. Their new device, which they called a transistor, was an electronic switch that took advantage of the fact that certain "semiconducting" crystals, such as silicon, sometimes conducted electricity and sometimes blocked it. By embedding impurities in a tiny fragment of semiconductor crystal, they found they could alter quite precisely how the crystal shifted between its conductive and insulating states in response to electrical currents. The transistor was therefore the simplest electrical switching device imaginable, with no moving parts; very small arrays of on-off switches are an ideal way to represent the 1s and 0s that are the language of a digital computer. It took a decade for companies to learn how to manufacture high-quality transistors in quantity; but once they were available, their small size, great durability, low rate of power consumption, and low rate of heat output quickly made them the technology of choice over vacuum tubes.

It was but a small step from the solid-state transistor to the integrated circuit—a single piece of silicon with two or more solid-state devices embedded in it and connected with wires made of thin layers of metal. Jack Kilby at Texas Instruments and Robert Noyce at Fairchild Semiconductor both independently invented the integrated circuit in 1959. The theoretical limits on the size of solid-state devices are vanishingly small, so as manufacturing technologies improved, integrated circuits could get more and more complex, squeezing enormous numbers of devices onto a single microchip of silicon. The smaller the devices, the less power they required, the less heat they produced, the lower the cost per device, and the faster and more powerful computers became. Noyce's and Kilby's first integrated circuits had only two devices on them, but within just a couple of years, even simple integrated circuits could reduce the size of a computer by 150 times. By 1970, state-of-the-art memory chips had 1,000 transistors; by 1980, the newest chips had 64,000 transistors; by 1992, 16 million transistors. A 1970 memory chip cost less than $10; a 1992 chip still cost less than $10.

IBM's entry into transistorized computing was smoothed by two massive 1950s contracts to provide computers for the SAGE and BMEWS early warning systems against Russian air and missile attacks. Defense computing made up a full half of IBM's computer revenues throughout the 1950s, and still funded 35 percent of its research in 1960. IBM's first fully transistorized computer, the 709, was built for the BMEWS project in 1958 and rolled out in a commercial version in 1960. High-speed ferrite core memories, magnetic tape drives, and flexible designs that could handle both business and scientific data requirements—all developed with government contracts—were recycled almost immediately into business products.

Other companies also won major defense contracts, but were typically much slower to adapt their products commercially. Rand's UNIVACs, for example, had to be painstakingly assembled on the customer's site, while IBM computers came in sleek, attractive casings, were easily installed, and were designed to perform operations that its customers really needed. The 650, a small business computer, sold more than a thousand units in the 1950s, becoming the world's first mass-produced computer. The 1401 became the industrial world's medium-sized mainframe standard at the end of the decade. The powerful 700 series dominated the high end of the market. Every new machine outsold the wildest marketing forecasts. The company's sales, which were only $40 million before the war, passed the billion-dollar mark in 1957. The CBS 1956 election night newscast pointedly used an IBM computer.

BETTING THE COMPANY

The decision to plunge ahead with the System/360 (later 360/370) mainframe computers was even more daring than the decision to go into electronic computing in the first place. The 1400 series of machines were hardly obsolete, and along with older IBM equipment, enjoyed the lion's share of installed back offices. But the fundamental objective of the System/360 was to replace virtually *all* other computers, including all the thousands of lease-paying IBM machines happily ticking away on customer sites (IBM equipment was almost always rented). As the 360 project leader, Vin Learson, wrote to a foot-dragging colleague, "Corporate policy . . . is that by 1967 the 1401 [IBM's then-current flagship] will be dead as a Dodo. Let's stop fighting this."

IBM's leading position was by no means beyond assault when it gambled on the 360. It had pioneered most of the business computer technology on the market, but now other companies—Philco, GE, RCA, Rand, Control Data—had learned from IBM's success. They were steadily chipping away at IBM's market share—cherry-picking opportunities where an IBM solution was becoming dated and vulnerable to newer products. The very breadth of IBM's success presented a daunting challenge—having created such a large and diverse new market, the dilemma was how to stay on top across the board.

The 360 decision was also forced by a drastic shift in the relative value of hardware and software in the short decade of commercial electronic computing. Only about 8 percent of the value of the earliest systems was accounted for by software. By the early 1960s, the software component was up to 40 percent of delivered value. Software development was becoming a major capital item for both IBM and its customers, and as systems proliferated, it was getting out of control. Customer reluctance to learn new software or rewrite their applications for each new generation of technology was looming as a major obstacle to continued market growth. To IBM, it was a crisis. Multiple machine designs and skyrocketing software costs pointed to an explosion of development expenses, since every important software product had to be redesigned for each class of machine.

The revolutionary new principle of the System/360 was *compatibility*, at a single stroke cutting through both the software problem and the breadth-of-market conundrum. Customers would be able to buy a range of computers, from a small $2,500/month machine up to an $115,000/month behemoth, a thousand times more powerful. But all the machines would run on the same software; better yet, IBM could emulate the 1400 software on the 360, so customers could, for the most part, transfer their 1400 programs directly to the 360. With the same software, the full range of machines would have the same "look" to the outside world, and so could plug into any of the full range of IBM memory units, printers, disk drives, and all the other peripheral equipment that supported a big computer operation. Before the 360, whenever a customer's operations outgrew its computer, the whole installation had to be ripped out, and a new system had to be put in almost from scratch—new equipment, new software, often new file designs. With the System/360, the installation could expand gracefully to meet the customer's needs—adding a faster processor here, more memory there, better software, upgraded printers, whatever. It was a much better deal for the customer, and also for IBM—as Watson put it, "Once a customer entered the circle of 360 users, we knew we could keep him there for a very long time."

It was a brilliant vision, but a huge gamble; the investment was so enormous that a botched product could have sunk the company, and for a very shaky time in 1965, it wasn't obvious that IBM would pull it off. IBM hired sixty thousand new employees, sank $750 million into engineering development, and opened five major new factories at a cost of $4.5 billion. The technical challenges were stupendous, and worse, were layered one on top of the other. The 360 was the first computer to use a hybrid integrated circuit—a way station to full integration—as a base technology. But no one had ever manufactured integrated circuits on the scale and at the quality level the 360 demanded. There was no recourse but for IBM, which had always bought almost all its electronic components, to create its own integrated circuit factories—at three times the cost of any previous computer factory—and create brand-new process technologies for the manufacture of integrated circuits on a mass scale.

Every peripheral component in the IBM product line had to be redesigned to assure the target compatibility throughout the series. More important, the software for the 360 series had to be consistent up and down the entire line. But that required millions of lines of code, the largest software program that had ever been written, and all under terrible time pressures. The cost overruns were appalling. At one point, when $600 million of inventory had been "lost," seemingly unfathomable metallurgical problems were shutting down the integrated circuit factories, and the huge software project was hopelessly bollixed, Watson admitted being close to panic, beset by fears that he had destroyed his father's company.

But the 360 pulled through. Not all the products were delivered on time, and the early versions underperformed specifications. Some of the software problems, despite a half billion dollars sunk into programming, bedeviled the system for years to come. But customers didn't mind waiting a bit. There was enormous loyalty to IBM and great confidence that it was a company that delivered on its promises. Even more important was the recognition that the 360/370 series was the *right* answer to a new era of pervasive high-performance computing. Once the kinks were out of the system and all the new factories were humming, IBM couldn't fill orders fast enough. The 360/370 series completely redefined the concept of modern business computing, and just as Watson predicted, once customers adopted the 360 concept, IBM owned them.

For twenty years thereafter, IBM's dominance of the industry was almost total. The formidable competitors of the 1950s and early 1960s—Burroughs, UNIVAC, NCR, Control Data, and RCA/Honeywell—were now just the "BUNCH." A new, and ignominious, sobriquet entered computer jargon—the "PCM," or Plug-Compatible Manufacturer. With no chance of taking on IBM frontally, competitors were reduced to manufacturing clone products, mostly peripheral devices and other equipment that could fit inside the 360/370 system architecture—like small animals darting in and out to snatch pieces of the lion's kill. By 1970, IBM's sales had soared to $7 billion, and kept growing, by more than 15 percent a year, every year, making it the largest and most profitable industrial company in history, with a grip on its industry that exceeded even that of Standard Oil or U.S. Steel in the turn-of-the-century heyday of unrestrained capitalist expansion.

IBM's success drew forth two new hostile forces, both more powerful than any single competitor. The first was its own home government, the second was the government of Japan.

UNITED STATES VS. IBM

Ramsey Clark, the attorney general of the United States during the last waning days of the Johnson administration, on January 17, 1969, signed a complaint charging IBM with unlawful monopolization of the computer industry and requested that the federal courts dismember the company. It was not the first time Justice Department's trustbusters had looked upon IBM's success with a jaundiced eye. When IBM's heavy investment in basic research began to pay dividends in the 1950s, the Justice Department initiated an action that forced IBM to license all of its patents, at a "reasonable" price to all comers, including the technology-hungry Japanese.

The 1969 suit had been expected for a long time, and IBM had already begun to unbundle the pricing of its systems, making it easier for other companies to sell compatible devices and software. But IBM's overall reaction to the case was arrogantly imperious. Fundamental antitrust theory, after all, in the words of Judge Learned Hand, held that a company was not to be penalized for its "superior skill, foresight, and industry." Surely, the courts, if not the Antitrust Division of the Justice Department, would quickly recognize that IBM had earned its market share simply by being better than everyone else.

In the event, the action dragged on for thirteen years. The government's case was meretricious from the outset, shot

through with contradictions and misconceptions. It took six years even to bring the case to trial, during which time the government repeatedly changed its theory of the case to keep pace with the dramatic changes in the industry and the constant turnover of Justice Department lawyers. The government's action brought forth a long series of private antitrust complaints. With the exception of a suit brought by Control Data before the antitrust action, which was settled out of court, IBM fought and won every one, twice on directed verdict—that is, without having to present a defense.

When the antitrust action finally came to trial, it was assigned to a federal judge, David M. Edelstein, whose behavior was frequently bizarre—at one point, he ordered depositions (previous sworn oral statements of witnesses) to be read aloud to an empty bench for seventy days. The case cost hundreds of millions of dollars, possibly as much as a billion, and spawned an entire industry of IBM-case lawyers and expert witnesses. As Frank Cary once put it, he told his legal staff to spend whatever was necessary, "and they still went over budget." IBM's chief expert witness, Professor Frank Fisher of MIT, named his new yacht *The Section 3* in honor of the key section of the antitrust statute.

An extensive review of the IBM case was an early priority of the Reagan administration Justice Department. The suit was dropped in June 1982 with a curt four-sentence appraisal by the solicitor general that the case was "without merit."

The government's action did great damage to the company. By the mid-1970s, the easy confidence of IBM executives that they would prevail on the merits had long since dissipated. The case had become a devouring monster, and the company was beginning to scrutinize every decision for how it might play in a hostile courtroom. Many of IBM's actions in the 1970s and 1980s, particularly its supine attitude toward small suppliers of PC components and software, can be explained as the reflexes ingrained by a decade in the courtroom's harsh glare. One of the more serious consequences of the antitrust case is that, in the anti-IBM atmosphere prevailing in the government, the United States typically refused to help against a much less principled assault being mounted by Japan.

Japan vs. IBM

IBM first established a presence in Japan in 1925 and reentered during the Occupation, so that as much as the postwar Japanese government might have liked, it could not banish IBM from its shores. By bargaining over import and manufacturing licenses, however, it gained extremely favorable royalty arrangements on the IBM patents and made them available to the entire Japanese industry. The American government took a generally benign view of such activities—it was happy to build up Japan as a democratic Asian power and equally happy to see someone cutting IBM down to size.

Japan's hardware manufacturing skills developed rapidly, and by the mid-1970s, Japanese research and production capabilities in certain semiconductor technologies were the equal of America's. But Japan had—and still has—a great weakness in software. Japanese companies could make good reproductions of IBM mainframes, but could not duplicate the IBM operating system, the extremely complex software that gave IBM computers their unique "look" to the outside world. Without the operating system, Japan's companies would be limited to making peripherals or other plug-ins for IBM environments and could not offer full substitutability for IBM systems.

The solution was massive, government-supported Japanese theft and industrial espionage against IBM. Throughout the 1970s, new IBM software, like the MVS operating system, which had taken years to develop, would show up in Japanese products almost immediately after its introduction. Japanese companies managed to fend off IBM's attempts to purchase their products, even through third parties, so IBM could not demonstrate the thefts. Fujitsu, the primary offender, has never sold mainframe operating system software in the United States, to avoid the reach of the federal courts and rules of discovery. Finally, in the early 1980s John Opel, IBM's chairman, refused to sell any further products or technology to Fujitsu unless they turned over their system software for inspection. When Fujitsu finally did so, the theft was obvious. A prolonged, multistage arbitration proceeding eventually awarded IBM $833 million in damages.

The climactic event in the history of Japanese thievery from IBM came in 1982. Bob O. Evans, who was director of IBM's advanced technology programs, received a telephone call in 1981 from a friend, Max Paley, a West Coast computer consultant and former IBM executive, asking what the "Adirondack Notebooks" were. The notebooks were an ultra-secret, eighteen-volume set of operating and software specifications for the next generation of IBM mainframes. Paley had been approached by Hitachi executives, who said they had half of the volumes, claimed to know precisely where the others were located, and hoped to hire Paley to obtain the notebooks and a Clark board, a key module in IBM's highest performance computers. (A trace on the stolen notebooks eventually disclosed that they had been removed and copied after an an IBM executive had left them for safekeeping in a Tokyo hotel safe. The hotel manager's brother was a Fujitsu employee. Fujitsu, in legendary "Japan, Inc." style, shared them with Hitachi, its bitter rival.) IBM, Paley, and the FBI mounted an elaborate sting operation that trapped Hitachi and Mitsubishi offering more than $600,000 in bribes to Paley, and suggesting an impressively large network of informers within IBM. In Evans's view, at least, the FBI terminated the operation prematurely for fear of losing their prey, but at the cost of not entrapping Fujitsu, who had stayed circumspectly on the fringes of the negotiations.

The initial reaction of the Japanese government and computer companies was outraged denial. There were demonstrations outside of IBM's Japanese headquarters and angry speeches in the Diet. IBM, however, finally got court approval to turn over three hundred hours of incriminating videotapes to the Japanese, bringing the protests to a shocked and embarrassed halt. Hitachi meekly agreed to pay IBM damages of $250 million and to allow IBM to inspect its new products for a period of five years. Incredibly, the Japanese government managed to swallow its embarrassment sufficiently to introduce a law into the Diet in 1984 to legalize software theft from foreigners—it was withdrawn only when IBM finally roused the American government to protest.

Dominating Information Technology

Reprehensible though it was, Japan's campaign to undercut IBM's dominant position in global computing demonstrated a much more sophisticated understanding of the industry than the American government's antitrust action.

Dominance in electronics is emphatically not the same as control. Standard Oil's dominance of the world petroleum industry at the turn of the century was based on its actual control of physical facilities—oil wells, pipelines, tankers, even barrel-stave forests. Computers and electronics are not that kind of industry. In the twenty years from the time IBM first established its leadership position in business computing, the price of processing power dropped by more than 100 times—processing power that cost $1 in the mid-1950s cost less than a penny in the 1970s. The price of computer memory dropped by more than 500 times, and the storage capacity of standard devices increased thousands of times. To maintain its lead position, IBM had to stay in front of an extremely fast-moving technology curve. The ratio of price to performance for the IBM 360 was about 50 percent superior to that of almost all competing machines. Instead of raising prices, in traditional monopolist style, IBM typically forced widespread price-cutting throughout the industry, always following up its initial offerings with a steady stream of new technology breakthroughs, such as faster memories, faster processors, and more extensive time-sharing capability. IBM's leadership was based not on controlling a technology, but on exploiting it better than anyone else.

By the mid-1970s, there were more than a hundred substantial companies competing in computers, most of which had not even existed when the 360 was announced. The 360's objective of compatibility over a broad product line—the ability to mix and match peripheral devices such as the central processor, memory, and input and output devices (like terminal keyboards, printers, and display screens)—required that the system be relatively "open." As long as a peripheral device conformed to specific electronic communication rules, it could be plugged in and work as a part of a seamless system. Openness and compatibility were opportunities for quick-moving companies that could offer products and software that were better and cheaper than IBM's own. Memorex and Storage Technology sold IBM-compatible memory equipment; Four Phase and Systems Engineering sold terminals; Texas Instruments, Intel, Ampex, and many other companies provided semiconductor components; Cullinane, Shared Medical Systems, and many others sold software. Data General, Prime, and Wyse created new classes of smaller, special-purpose computers that would work within an IBM computing environment. By the middle of the 1970s, Amdahl, part-owned by Fujitsu, but the creation of Gene Amdahl, one of the 360's key designers, was producing IBM mainframe clones, completely look-alike machines that matched, or some customers said, exceeded, IBM's performance point for point.

Clearly, IBM did not control the computer industry in any classic sense, and as competitors proliferated, and governments forced broader and broader licensing, its overall market share inevitably declined. But, paradoxically, IBM's sales and profitability continued to bound ahead, and its dominance over the industry became, if anything, even more pervasive. By the end of the 1970s, 70 percent of the world's computer installations were still centered around IBM equipment, and IBM's share of world computer profits was even higher.

The secret of IBM's dominance, as IBM itself understood better than anyone, was that it had created, and owned, a pervasive industry *architecture.* All the competitors were playing by IBM's rules—making devices, writing software, manufacturing clones, running time-share centers—all within a computing environment that IBM defined and that only IBM had completely mastered. The confidence of IBM customers was so great, their commitment to the 360/370 architecture so deep, that no competitor had a chance of replacing it. It would mean throwing out too much investment built up over too long a time.

The consequence was that no one could beat IBM to market with a new product line. If a competitor tried to invade its space *ahead* of IBM, it could never be sure that IBM's next operating system release would be compatible with its product, especially if the product was one IBM wanted for itself. Competitors had no choice but to reverse-engineer IBM products only *after* they became available, and therefore were condemned always to be second to market. And by the time competitive plug-compatible products became available, IBM was usually already moving on to the next product generation. That explains the desperation of the Japanese to steal information in advance of its release.

IBM's "lock" on the mainframe industry was never complete and was under constant challenge from almost every quarter. The astonishing pace of technological advancement in electronics demanded continual leapfrogging innovation to stay apace. The loyalty of IBM's customer base depended on a deeply engrained confidence that, on average, IBM technology would always be the best, that its price/performance would always be at the top, that the industry-shaking innovations would always come from IBM first. It was an extraordinarily demanding game, and IBM played it brilliantly for more than twenty years, an accomplishment that ranks as one of the signal industrial achievements in all of business history. It was easy to poke fun at Big Blue—the stiff uprightness, the formalities of the blue suits and white shirts, the flat inflections of computerspeak. But IBMers knew, and they knew the rest of the world knew, that IBM was truly the best at what it did, outdistancing its competition to a degree that perhaps no company ever had.

So it is the more surprising that IBM's top management missed the straws swirling in the wind in the 1970s and 1980s that computing was undergoing fundamental change, that the 360/370—indeed, the whole mainframe principle—was heading for a dead end. But one of the most important of those straws had been floated by IBM itself, who had already, by 1982, made the personal computer the product of the decade.