Time Machine

A Brief History of Computing

Graeme Philipson

This is the original text of a long article that appeared in the December 2003 edition of Australian Personal Computer. The published version differs a little from this.

Time Machine – A Brief History of Computing

The Urge to Compute

In 1901 divers exploring a sunken ship off the Greek island of Antikythera made a remarkable discovery. Among the statues and urns in the pre-Christian wreck they found a sophisticated mechanism consisting of a number of toothed wheels cut from sheets of bronze, which meshed together to perform a vast range of astronomical calculations.

Nothing like the Antikythera mechanism had ever been discovered. Nor has anything since. Dating techniques give it a vintage of around 100 BC, an era when it was not previously thought mankind was capable of making such devices. It displays a level of engineering that did not become widespread until the industrial revolution.

The Antikythera has become the most famous of early calculating devices. It is also the oldest that has been found, but it is not alone. Since man first began counting, he has used mechanical devices to help him add and multiply. The Greeks and Romans used counting boards, often as an aid to gambling, which evolved into the abacus. Early Christian missionaries took the abacus to China, where it reached its highest levels of sophistication.

Meanwhile numbering systems evolved to allow more complex calculations. The biggest advance was the invention of zero by Indian mathematicians in the ninth century. What we call Arabic numbers are in fact Hindu, and were bought to Europe by the Arabs in the middle ages ("algebra" is an Arabic word).

Modern computing began in the 17th century with John Napier, a minor Scottish nobleman. Napier made two significant mathematical advances, both of them important in the history of computing. The first was the calculating aid known as Napier's Bones, rods or cylinders which slid together to mechanise the multiplication tables. The second was his invention of logarithms, which employed the mathematical technique of numerical powers to simplify the difficult tasks multiplication and division into the much easier jobs of addition and subtraction.

Log tables were a staple of mathematical calculation until very recently. Their use was taught in Australian high schools until the 1970s, when electronic calculators rendered logs and their mechanical cousins, slide rules, obsolete. Logarithms influenced generations of mathematicians, and led directly to the development of the first mechanical calculators. Shortly after Napier died in 1617, three remarkable men independently developed the first working mechanical calculators.

The first was Wilhelm Schickard, whose "Calculating Clock" was devised in 1623 to help astronomer Johann Kepler with his celestial calculations. Schickard's machine perished in the Thirty Years War, and its existence was not rediscovered by history until his papers were found in the 1930s.

The second was the "Pascaline", invented by the famous French mathematician Blaise Pascal in 1642. The Pascaline was an adding machine that used toothed wheels to transfer the digits by the "nines complement" method. Many variants were constructed, and the device became widely used and very well known, but it had many limitations. It was very complex, and prone to mechanical failure.

The most sophisticated 17^{th} century device was made by the German polymath Gottfried Wilhelm von Leibniz, the man who invented differential calculus in the form we know it today. (Isaac Newton independently invented a more cumbersome version a little earlier). Leibniz was also the first man to study at any length binary arithmetic, which to his mind proved that God had made the world out of nothing.

Liebniz's remarkable "Stepped Reckoner" was a major advance on the Pascaline. It employed a "Leibniz wheel", a long cylinder that ran the length of the device and enabled multiplication as well as addition. But, as often happened with devices from this era, Leibniz's conception was greater than the technology's ability to build devices of such sophistication, and the machine could rarely be made to actually work. That problem was also to plague the inventor of the most sophisticated mechanical device, and the first true computer, Charles Babbage.

The Father of Computing

Charles Babbage was born in London in 1791, the son of a wealthy banker. He was a brilliant mathematician and one of the most original thinkers of his day. His privileged background gave him the means to pursue his obsession, mechanical devices to take the drudgery out of mathematical computation. His last and most magnificent idea, the Analytical Engine, can lay claim to being the world's first computer.

By the time of Babbage's birth, mechanical calculators were in common use throughout the world. Other inventors had built on Pascal's and Leibniz's work, and manufacturing techniques had improved to the extent that their designs could be made to work. But they were calculators, not computers – they could not be programmed.

Nor could Babbage's first conception, which he called the Difference Engine. This remarkable device was designed to produce mathematical tables. It was based on the principle that any differential equation can be reduced to a set of differences between certain numbers, which could in turn be reproduced by mechanical means. The story has it that Babbage and his astronomer friend John Herschel (son of the more famous Sir William Herschel, who discovered Neptune) were lamenting the many inaccuracies in the mathematical tables of the day. Babbage remarked that he wished the tables could be produced by steam engines. Herschel said he believed they could, so Babbage set about pursuing the idea.

The Difference Engine was the result. It was a far more complex machine than anything previously conceived. Partially funded by the government, and partially by his sizeable inheritance, Babbage laboured on it for nearly twenty years, constantly coming up against technical problems. As had been the case with Pascal and Leibniz, the device was more complex than what could be made by the machine tools of the day. But he persevered, and was eventually able to construct a small piece of it that worked perfectly and could solve second-level differential equations.

The whole machine, had it been completed, would have weighed two tonnes and been able to solve differential equations to the sixth level. After battling with money problems, a major dispute with his grasping chief engineer, the death of his wife and two sons, and arguments with the government, the whole project collapsed. Part of the problem was Babbage's perfectionism – he revised the design again and again in a quest to get it absolutely right.

By the time he had nearly done so he had lost interest. He had a far grander idea – the Analytical Engine – which never came close to being built. This remarkable device, which lives on mainly in thousands of pages of sketches and notes that Babbage made in his later years, was designed to solve any mathematical problem, not just differential equations.

The Analytical Engine was an horrendously complex device containing dozens of rods and hundreds of wheels. It contained a mill and a barrel, and an ingress axle and egress axle. Each of these components bears some relationship to the parts of a modern computer. And, most importantly, it could be programmed, by the use of punched cards, an idea Babbage got from the Jacquard loom.

But none of it was to be. Babbage died in 1871, his life's work unfinished. His work was widely known in his lifetime, but only in recent years has the breadth of his vision been fully appreciated. Australian academic Allan Bromley, before his death in 2001, became the world's leading authority on Babbage. Bromley spent years poring over Babbage's notes, demonstrating beyond doubt that Babbage's designs represent the world's first computer, if only in conception.

In 1995, largely on Bromley's advice, Sydney's Powerhouse Museum paid $400,000 for a large piece Charles Babbage's famous Difference Engine Number 1, one of only three such left in the world. It is permanently on display, the showpiece of the museum's impressive collection of computing artefacts. Babbage today has the recognition he deserves, as the Father of Computing.

The first programmer

Ada, Countess of Lovelace, was daughter of the famously dissolute English poet Lord Byron. She can also lay claim to being history's first computer programmer.

Augusta Ada Byron was born in 1815 and brought up by her mother, who threw Byron out, disgusted by his philandering. She was named after Byron's half sister, who had also been his mistress. Her mother was terrified that she would become a poet like her father, so Ada was schooled in mathematics, which was very unusual for a woman in that era. But she did not entirely forget her heritage – her understanding of mathematics, and they way she described the subject, was very lyrical.

She met Charles Babbage in 1833 and became fascinated with the man and his work. In 1843 she translated from the French a summary of Babbage's ideas which had been written by Luigi Federico Manabrea, an Italian mathematician. At Babbage's request she wrote some "notes" that ended up being three times longer than Manabrea's original.

Ada's notes make for fascinating reading. "The distinctive characteristic of the Analytical Engine … is the introduction into it of the principle which Jacquard devised for regulating, by means of punched cards, the most complicated patterns in the fabrication of brocaded stuffs … we may say most aptly that the Analytical Engine weaves algebraical patterns just as the Jacquard-loom weaves flowers and leaves."

Her notes included a way for the Analytical Engine to calculate Bernoulli numbers. That description is now regarded as the world's first computer program.

Ada Lovelace's life was beset by scandalous love affairs, gambling and heavy drinking. Despite her mother's best efforts, she was very much her father's daughter. She considered writing a treatise on the effects of wine and opium, based on her own experiences. She died of cancer in 1852, aged only 37.

Herman Hollerith and the Birth of IBM

The 1880 US census counted 50,155,783 inhabitants in the republic. The problem is, the task of tabulating the data was so massive that it was not completed until 1888, just two years before the next census was due. Clearly, something had to be done.

The new census superintendent, an Englishman named Robert Porter, announced a competition to develop a new system that would enable the results of the 1890 census to be returned more quickly. The winner, by a large margin, was a system developed by a young engineer called Herman Hollerith.

Hollerith's method used punched cards to store information. Though it was no quicker than its rivals in transcribing census information, it was ten times faster in tabulating the results. Every person's details, such as age, sex, and ethnic origin, was stored on a single punched card, and these cards could then be sorted and counted automatically.

This was achieved through the employment of an electromechanical device called a tabulator, which sensed the absence or presence of a hole in the card in a particular position, then added the totals. Just six weeks after counting began, the Census was able to announce that the USA's population had increased to 62,622,250 people in ten years. All cross-tabulations were completed in just two and a half years, compared to over seven years for the previous census.

Hollerith's machine was a sensation, and it was quickly adopted by other census departments around the world. In modified forms it was also widely employed by other organisations with large data collection and storage problems, such as banks and insurance companies. It spawned many imitators, and by World War I punched cards became a way of life in the business world.

This was an era of unprecedented growth and of rapid technological advance on many fronts. Electricity, aircraft, automobiles and radio all date from this era. So do typewriters, cash registers and adding machines. By the 1920s a large office machinery industry had grown up, most of it based in the USA, which emerged from World War I as the world's most prosperous and dynamic economy.

In 1911 Hollerith merged his firm, the Tabulating Machine Company, with two others to form the Computing-Tabulating-Recording company, known as C-T-R. In 1914 the company hired a new general manager, Thomas J. Watson, who had been the number two man at National Cash Register (NCR), the leading cash register company. Ten years later, after he became president, Watson changed the company's name to International Business Machines, better known by its initials IBM.

Watson learnt most of his business techniques from John Patterson, the legendary founder of NCR. Patterson demanded total loyalty from his employees, down to the clothes they wore and the sports they played. His anti-competitive business practices, for which he was indicted in 1914, were also widely copied by IBM, and have echoes in IBM's antitrust suit in the 1970s and Microsoft's behaviour today. The computer industry has a very long history.

A new kind of maths

George Boole (1815-1864) was a self-taught English mathematician who made the most significant advances in mathematical logic since the ancient Greeks. He invented what has come to be known as Boolean algebra, a kind of binary logic that is the basis for the way modern computers process instructions and data.

At the heart of Boolean algebra are instructions like AND, NOT and OR, logical operators that explain the behaviour of sets of numbers when they interact with each other. In his seminal books "The Mathematical Analysis of Logic" (1847) and "Investigation of the Laws of Thought" (1854), he combined mathematics and logic into a unified system without which modern electronic computation would be impossible.

Boole's work was extremely influential on the subsequent development of both mathematics and formal logic. His work can now be seen as a necessary precondition to the development of computers, and in particular the later work of Turing and von Neumann.

Calculation goes to War

On 3 September 1939 German tanks rolled into Poland. World War II had begun. Even as the Panzers attacked, the cream of Poland's cryptanalysts were packing up their papers and leaving for England and France.

Over the previous ten years these unsung Poles had performed one of the greatest feats in the history of cryptanalysis. They had succeeded in reverse engineering the German Enigma machine, the famous German mechanical encryption device.

The Enigma machine was first developed by the small German army, the Reichswehr, in the 1920s. After Hitler came to power in 1933 Germany quickly rearmed, and the Enigma program was vastly expanded. Hundreds of machines were built, of ever greater complexity, and they formed the basis of all military and diplomatic communications.

The Enigma consisted of a set of rotors – three at first, but later expanded to many more – that could set into a prearranged pattern. Any similar machine that had its rotors set in the same pattern could decode any message sent by the first. It was a variant of the standard substitution method of cryptography used throughout time, but made vastly more difficult to decode by the use of mechanical keys.

Difficult, but not impossible. During the 1930s Poland's cryptographic department were able to decode Enigma's messages by simulating its processes in an electrical device called a

Bombe. After Poland fell, this expertise was transferred to the British decoding facility at Bletchley Park, which found itself faced with ever more complex versions of the Enigma.

The British were faced with much bigger problems than the Poles. Every new rotor the Germans added to the Enigma meant a thousandfold increase in the complexity of the key. The Polish Bombes were unable to handle the increased calculating load. The answer was the Colossus, an electronic decoding machine with 2400 vacuum tubes. Ten were built, enabling the British to electronically scan Enigma messages at 25,000 characters a second.

Heading the British efforts at Bletchley Park was Alan Turing, an eccentric mathematician who had before the war written a paper on "computable numbers", in which the concept now know as a Turing Machine was born. Turing proposed a simple theoretical device that could scan a tape with symbols on it, then read, write or erase each symbol on that tape as it moved left or right.

This was a development in symbolic logic as significant as George Boole's, and as just important in the development of computing. Very simple in concept, the Turing Machine describes how modern computers read, process and write data. Turing remains one of the most significant figures in computing history, his life cut short by suicide in 1954 after he was exposed as a homosexual in a society with little tolerance for such behaviour.

The most significant developments during World War II came from the British, but it could all have been very different. Konrad Zuse was a brilliant German engineer who single-handedly developed the most complex devices in the annals of the early history of the computer industry, but whose work was ignored by the Nazis and ended up in an historical dead end.

Zuse built four machines before and during the war. Most remarkable was that he worked almost totally alone, with no knowledge of what was happening elsewhere. His first machine, the Z1, was entirely mechanical, but used binary arithmetic, remarkable in itself. The Z2 was electromechanical improvement, and in 1942 he proposed an improved general purpose computer to the Wehrmacht. He was turned down, but he persuaded the Aerodynamics Research Institute to fund the Z3, which is now on display in the Deutsches Museum in Munich.

These devices were not programmable, and were therefore not true computers, but they used techniques and technologies that are essential components of modern computer systems. But for the fortunes of war Zuse, who died only in 1995, would today be remembered as the father of electronic computing.

ENIAC - The First True Computer

While the British were decoding the Enigma machines and Konrad Zuse was working alone in Germany, even more significant developments were occurring across the Atlantic. Scientists at Bell Labs, MIT, Harvard and elsewhere were all building electromechanical calculators of various sorts.

Their efforts were intensified when the USA entered the war after Pearl Harbour. IBM and Harvard's Howard Aiken developed a giant electromechanical calculator called the Mark I, which was essentially a realisation of Babbage's dream nearly a century after the act. But it was a dead end – the future lay in electronics.

The first true electronic computer was the ENIAC, which stood for Electronic Numerator, Integrator, Analyzer and Computer. ENIAC dates from 1942, when a 35 year old engineer named John W Mauchly wrote a memo to the US government outlining his ideas for an "electronic computor". His ideas were ignored at first, but they were soon taken up with alacrity, for they promised to solve one of the military's most pressing problems.

That was the calculation of ballistics tables, which were needed in enormous quantities to help the artillery fire their weapons at the right angles. The US government's Ballistics Research Laboratory commissioned a project based on Mauchly's proposal in June 1943. Mauchly led a team of engineers, including a young graduate student called J. Presper Eckert, in the construction of a general purpose computer that could solve any ballistics problem and provide the reams of tables demanded by the military.

The machine was to use vacuum tubes, a development inspired by Mauchly's contacts with John Atanasoff, who used them as switches instead of mechanical relays in a device he had built in the early 1940s. Atanosoff's machine, the ABC, was the first fully electronic calculator. Mauchly's decision to follow his lead and use vacuum tubes was to prove very significant.

ENIAC differed significantly from all devices that went before it, except Zuse's unnoticed efforts in Germany. It was to be programmable. Its use of stored memory and electronic components, and the decision to make it a general purpose device, mark it as the first true electronic computer.

But despite Mauchly and Eckert's best efforts ENIAC, with 17,000 vacuum tubes and weighing over 30 tonnes, was not completed before the end of the war. It ran its first program in November 1945, and proved its worth almost immediately in running some of the first calculations in the development of the H-Bomb (a later version, appropriately named MANIAC, was used exclusively for that purpose).

ENIAC led to EDVAC (Electronic Discrete Variable Computer), incorporating many of the ideas of John von Neumann, a well-known and respected mathematician who lent a significant amount of credibility to the project. Neumann also bought significant intellectual rigour to the team, and his famous paper "report on EDVAC" properly outlined for the first time exactly what an electronic computer was and how it should work.

The war over, Mauchly and Eckert decided to commercialise their invention. They developed a machine called the UNIVAC (Universal Automatic Computer), designed for general purpose business use. But they were better engineers than they were businessmen, and after many false starts their small company was bought by office machine giant Remington Rand in 1950. The commercial computer industry had begun.

Australia's first computer, and one of the first in the world, was CSIRAC. It was built by Trevor Pearcey for the CSIRO and commissioned in 1949. It is on display in the Museum of Victoria, and though it no longer works it is the oldest computer still in existence. Pearcey, who died in 1998, was a true computing pioneer, and learnt his craft in the UK during and after World War II. But his efforts were largely unappreciated, and after the Menzies government's directions that CSIRO should concentrate on agricultural science Australia became the IT colony it remains today.

The Birth of the Computer Industry

Computers had arrived. ENIAC, EDVAC, CSIRIC, UNIVAC, ILLIAC, SILLIAC, MANIAC, BINAC. But the early computers were all used for scientific means, and it was 1951 before the first commercial machine was delivered. It was a UNIVAC, installed in the US Census.

UNIVAC leapt to the forefront of public consciousness in the 1952 US presidential election, where it correctly predicted the results of the election based on just one hour's counting. It was not a particularly impressive machine by today's standards (it still used decimal arithmetic, for a start), but nearly 50 of the original model were sold.

The 1950s was a decade of significant improvements in computing technology. The efforts of the Bletchley Park crew during the war led to a burgeoning English computer industry. Before his death, after studying von Neumann's EDVAC paper, Alan Turing designed the ACE (Automatic Computing Engine), which led to the Manchester Mark I, technically a far superior machine to ENIAC or EDVAC. It was commercialised by Ferranti, one of the companies that was later to merge to form ICL, the flagbearer of the British computer industry.

The most significant US developments of the 1950s were the Whirlwind and SAGE projects. MIT's Whirlwind was smaller than ENIAC, but it introduced the concepts of real-time computing and magnetic core memory. It was built by a team lead by Ken Olsen, who later founded Digital Equipment Corporation, the company that led the minicomputer revolution of the 1970s.

SAGE was a real-time air defence system built for the US government in the Cold War. The project was accorded top priority, with a virtually unlimited budget. In a momentous decision, the government awarded the contract to a company that had only just decided to enter the computer industry. That company's name was IBM.

SAGE (Semi-Automatic Ground Environment) broke new ground on a number of fronts. The first was its sheer size. There were 26 data centres, each with a 250 tonne SAGE mainframe. It was built from a number of modules that could be swapped in and out. It was the world's first computer network, using the world's first fault-tolerant computers and the world's first graphical displays. And it gave IBM a head start in the computer industry that it has retained ever since.

By the end of the 1950s there were dozens of players in the computer industry. Remington Rand had become Sperry Rand, and others like RCA, Honeywell, General Electric, Control Data and Burroughs had entered the field. The UK saw the likes of Ferranti and International Computers and Singer, and continental Europe Bull and Siemens and Olivetti. In Japan, a 40 year old company called Fujitsu moved into computers.

Every computer from each of these suppliers ran different software. They used different printers, different storage devices, and stored their data in different ways. IBM decided to change all that. In the single most momentous development in the history of computing, IBM in 1964 released the System/360, its revolutionary family of mainframes.

A team of engineers, led by the legendary Gene Amdahl, designed a modular computer system that could be upgraded by the addition of more memory or bigger disks, and which ran the same operating system from the bottom to the top of the range. This seems normal now, but at the time IBM's approach was revolutionary. It meant a massive investment, one that was so large that IBM's senior management later admitted that it was a "bet the company" gamble. And it very nearly didn't come off – the sheer size of the project meant severe time and cost overruns.

The S/360 gamble (the numbers were meant to indicate the points of the compass) paid off handsomely. The S/360 evolved into the S/370, and then into the S/390. It is still in existence today, and its success meant that every other computer manufacturer subsequently followed the same architectural approach. The S/360 marked the beginning of the modern computer industry.

Minis and Micros

By the mid 1960s computers were in common use in government and industry throughout the world. Owing to the tremendous success of the S/360, IBM was the industry leader, as large as its major competitors combined.

Those competitors were collectively called the "Bunch" – a clever acronym for Burroughs, Univac (by this time the company was called Sperry, though it still used Univac as a model name), NCR, Control Data and Honeywell. IBM and the Bunch sold large and expensive mainframe computers to government departments and large corporations. But some people thought that computers need not be that big, nor cost that much.

One of these was Ken Olsen, who had worked on the Whirlwind project. In 1957 he started a small company called Digital Equipment Corporation, better known as DEC, in an old wool mill in Boston. DEC's first products were small transistorised modules that could be used to take the place of the vacuum tubes still used by computers of that era.

These modules proved very popular, and DEC was soon making so many different types that Olsen decided to build his own computers based around them. The first of these, the PDP-1 (the PDP stood for Programmed Data Processor), was released in 1960. It was followed by the PDP-5 in 1963 and the PDP-8 in 1965.

The PDP-8 was revolutionary. It ushered in the minicomputer revolution and brought computing to a whole new class of users. It had just 4K of memory, but it was a real computer, and much less expensive ($US18,000) than any other machine on the market. It was an enormous success, especially with scientists and engineers, who finally had a computer they could afford and easily use.

In 1970 DEC released the equally successful PDP-11. New technology and economies of scale meant that the prices of DEC's minicomputers kept dropping as quickly as their capabilities improved, and soon DEC had many competitors. One of the most successful, Data General, was founded in 1968 by ex-DEC employees. The Data General Nova, announced in

1969, set new benchmarks for price-performance, and by 1970 over 50 companies were making minicomputers.

These included IBM and the Bunch, but they moved slowly and largely missed the act. Only IBM caught up, by the end of the 1980s, with its AS/400. The big winners were Data General and other startups and companies new to the industry like Prime, Hewlett-Packard and Wang. In 1977 DEC released the VAX, the world's most successful minicomputer, still in use today. DEC was acquired by Compaq in 1997, then became part of Hewlett-Packard in 2002.

If smaller was better, smaller still was better still. Some people even believed that computers could be made small enough and cheap enough that they could be bought and used by individuals. For this to happen, computers needed to get much smaller and cheaper. Technology, as always, was to make this possible.

Early computers used vacuum tubes as switches. These were soon replaced by transistors, invented by Bell Labs' William Shockley, John Bardeen and Walter Brattain in 1948. The three received the Novel prize for their achievement. Transistors worked in solid state – the switching depended on the electrical properties of a piece of crystal – and led to further developments in miniaturisation throughout the 1950s and 1960s. The next significant development was the invention of the integrated circuit (IC) by Jack Kilby at Texas Instruments in 1959. ICs combined many transistors onto a single chip of silicon, enabling computer memory, logic circuits and other components to be greatly reduced in size.

Electronics was becoming a big industry. After he left Bell Labs, Shockley started a company to commercialise the transistor. His acerbic personality led to problems with his staff, and eight of them left in 1957 to found Fairchild Semiconductor. Two of those eight, Gordon Moore and Bob Noyce, in turn founded their own company, Intel, in 1968. All these companies were located in the area just north of San Jose, California, and area dubbed "Silicon Valley" in a 1971 Electronic News article by journalist Don Hoefler.

The Microprocessor

The next step was the development of the microprocessor, conceived by Intel engineer Marcian Hoff in 1970. Hoff's idea was simple. By putting a few logic circuits onto a single chip, the chip could be programmed to perform different tasks. The first microprocessor, the 4004, was developed as a cheap way of making general purpose calculators.

The 4004 was not an immediate success. No-one really knew what to do with, but sales picked up as its flexibility became apparent. In 1972 Intel released the 8008, and then in 1974 an improved version called the 8080. People started to realise that these devices were powerful enough to run small computers.

In July 1974 Radio-Electronics magazine announced the Mark 8, "Your Personal Minicomputer", designed around the 8008 by Virginia postgraduate student Jonathan Titus. You had to send away for instructions on how to build it, but thousands did. It success inspired rival magazine Popular Electronics to announce the "World's First Minicomputer Kit to Rival Commercial Models" in its January 1975 issue. The device, the Altair 8800, was designed by Ed Roberts, who ran a small electronic company in Albuquerque called MITS.

MITS sold Altair kits for less than $US400, at a time when the cheapest DEC PDP-8 cost more than ten times as much. Roberts was swamped with orders, and the microcomputer revolution had begun.

One of the Altair's great strengths was its open architecture. Roberts deliberately designed it so that others could add to it by developing plug-in cards. As sold by MITS, the device was very limited, but hobbyists and small companies all over America soon began developing hardware and software for it.

This included two Harvard undergraduates called Paul Allen and Bill Gates. Both had been playing with computers since their high school days in Seattle. Allen noticed the Altair story in the magazine, and suggested to Gates that they write a BASIC interpreter for it. BASIC was a simple language that enabled computers to be more easily programmed – the only way to use the first Altair was by setting switches.

Allen rang Roberts in Albuquerque, Gates wrote the compiler in six weeks, and the two drove to New Mexico, Gates finishing off the software in the parking lot before their meeting with Roberts. The compiler worked, and Gates and Allen dropped out of Harvard and started a company they called Micro-Soft around the corner from MITS in Albuquerque. Soon after they dropped the hyphen and moved their small company back to their hometown in the Pacific Northwest.

The Altair spawned a host of imitators. Computer clubs sprang up across the world. The most famous, located of course in Silicon Valley, was the Homebrew Computer Club. Two of the club's most active members were Steve Wozniak and Steve Jobs, who teamed up to build a little computer called the Apple I, powered by the 6502 microprocessor from a small Silicon Valley company called MOS.

The Apple I was moderately successful, so the two Steves decided to go into business properly. Jobs sold his VW microbus and Wozniak his HP calculator, they borrowed $5000 off a friend, and they were in business. They soon attracted the attention of venture capitalist Mike Markkula, who believed that microcomputers were the Next Big Thing. He was right.

Apple released the Apple II in the middle of 1977, about the same time as commercial designs from Tandy (the TRS-80) and Commodore (the PET). But the Apple II outsold them both, because of its attractive design and its greater ease of use. In 1980 Apple went public, in the most successful float in Wall Street history. That year Time magazine named the microcomputer "Man of the Year.

Yet the revolution had hardly begun.

The IBM PC

The success of the Apple II and other early microcomputers persuaded IBM to enter the market. In July 1980 Bill Lowe, head of IBM's entry level systems division made a presentation to IBM senior management about why Big Blue should make a move. More importantly, he suggested how this could be done.

The key to doing it quickly, said Lowe, was to use standard components. This was a major departure for IBM, which normally designed and built everything itself. There was no time for that, argued Lowe. Management agreed, and he was told to go off and do it. The building of the IBM PC was given the name Project Chess, and the machine itself was internally called the Acorn.

The machine was ready in less than a year. It was a triumph of outsourcing. The microprocessor was an Intel 8088. Microsoft supplied the operating system (see box) and a version of the BASIC programming language. Disk drives (just two low capacity floppies) were from Tandon, printers from Epson, and power supplies from Zenith. Software included a word processor and a spreadsheet.

Many in IBM were uncomfortable with the idea that the company should become involved in the personal computer market. One famous internal memo warned that it would be “an embarrassment” to IBM. The doubters were quickly proved wrong. Within days of the machine’s launch on 12 August 1981, IBM was forced to quadruple production.

Still they could not keep up with demand. People, and especially businesses, who were previously wary of microcomputers were reassured by the IBM logo. A brilliant advertising campaign featuring a Charlie Chaplin lookalike hit just the right balance between quirkiness and quality. The machine was no technological marvel, but it worked, and of course it was from IBM.

Big Blue’s decision to source the components from other manufacturers had far-reaching, if unintended, consequences. It meant that anybody could copy the design. Hundreds of companies did, and the IBM PC became the industry standard. A huge industry grew up in peripherals and software, and for the first few years the big battle in the computer industry was over “degrees of compatibility” with IBM.

Some companies, it seemed, were even more IBM-compatible than IBM, extending the architecture even further. One such company, Compaq, grew out of a 1982 meeting in a Houston pie shop when Rod Canion, a senior manager with Texas Instruments, sketched a picture of a portable IBM-like computer on a napkin. The first Compaq, portable like a sewing machine is portable, hit the market just a few months later, and by 1985 the company was in the Fortune 500. It was the fastest any company had ever got there.

At the beginning of the 1980s virtually no-one used microcomputers. By the end of the decade they were on every desk. The growth in the PC industry during those ten years remains one of the most remarkable tales in the history of business, and in the history of technology.

Like all the other components of the first PC, IBM bought the operating system off the shelf. It initially called on a company called Digital Research, developers of the CP/M operating system used on many early microcomputers. But Gary Kildall, Digital Research’s idiosyncratic founder, broke the appointment because he was out flying his plane. Irritated, IBM turned to another small company they had heard had a suitable operating system. That company’s name was Microsoft.

The problem was, Microsoft did not have an operating system. Bill Gates didn’t let IBM know that. He quickly bought an operating system called QDOS (which stood for “quick and dirty operating system") from another small company, Seattle Computer Products, for $US30,000. He renamed it MS-DOS and licensed it to IBM. The licensing deal was important – for every IBM PC sold, Microsoft would receive $US40. Microsoft previously just another minor software company, was on its way.

The Challenge from Apple

During most of the 1980s, it seemed like Apple would fulfil its early promise, and perhaps become the dominant force in the PC industry. It answered the challenge from IBM's PC with the Macintosh, an innovative machine that brought the graphical user interface (GUI) to the people. The GUI had been invented at Xerox PARC (see box, p93) years earlier, and Apple had already used it on its earlier ill-fated Lisa computer. But the Mac made it work.

The Macintosh was announced with a single advertisement during the 1984 Superbowl. The ad pictured a woman running at a screen with sledgehammer, mounting the stage, and shattering the image of a Big Brother type figure indoctrinating the faceless masses. It was a sensation.

So was the Mac. It was underpowered, and it lacked software, but these shortcomings were soon addressed and the Mac looked like it would live up to its early slogan – "the computer for the rest of us." Apple developed a low-cost laser printer – also invented at Xerox PARC – and the first piece of publishing software, and overnight the desktop revolution was born.

Within eighteen months or so the publishing industry worldwide was stood on its head. Typesetters and printers went out of business by their thousands. Now anybody could publish professional quality newsletters, posters, and magazines for a fraction of the cost of a few years earlier.

But Apple blew it. Unlike the IBM PC, the Macintosh had a closed architecture, which Apple protected closely. It refused to license its technology, which meant there were no direct competitors, and Apple's prices remained high. That did not matter at first, because the technology was so superior, but gradually the PC caught up. When Microsoft released a usable version of Windows in 1993, the Mac's technical superiority all but disappeared.

Apple did eventually license its technology, and Apple clones began appearing in the late 1980s. But, in one of the greatest acts of corporate lunacy ever witnessed, Apple reversed its decision, sealing its fate as a niche player with little advantage over its competitors except a nice logo. It came within an ace of merging with IBM in 1994 (see box p100), but when that fell through Apple was left with few friends beyond its declining band of loyal users.

At the end of the 1980s it was far from certain which computer operating system, and which hardware architecture, would win out. Apple was still strong, and in the IBM PC and compatible world there was substantial uncertainty over which operating system would win out. MS-DOS was still dominant, but it faced challenges from Digital Research's DR-DOS and – more significantly – from IBM's OS/2.

After the PC's astounding success, IBM realised it had made a mistake licensing an operating system off Microsoft. The operating system was where the battle for the hearts and minds of users was being won, so it determined to wrest control back from the uppity boys in Seattle. OS/2 was the answer.

OS/2 was a vastly superior operating system to MS-DOS. It had a microkernel, which meant it was much better at multitasking, and it was built for the new era of microprocessors that were then coming out. It would operate across architectures and across platforms. In its early days, Microsoft and IBM cooperated on its development, but Microsoft withdrew to concentrate on something it called Windows NT, which stood for New Technology.

We all know what happened next. Microsoft outmarketed IBM, and OS/2 died. Apple imploded, Microsoft started bundling its applications, and the battle for the desktop became a one-horse race. The scene was set for the industry's next big shakeup, the rise of the Internet.

The Big Blue Apple?

The biggest "if" in the history of computing came within an hour of actually happening. In November 1994 senior management of Apple and IBM met in a hotel room in Chicago to thrash out the final details of a merger between the two companies.

The meeting was the culmination of months of tough negotiations. IBM and Apple had forged a technology alliance in 1991, and key figures in both companies saw major advantages in joining the two companies closer together. IBM had the marketing clout, the cash, and the high-end business market, and Apple had the technology, the flair, and important parts of the microcomputer market. To backers inside both companies, it was a deal made in heaven.

The meeting began at 8:30am. In attendance were Apple's CEO Michael Spindler and chairman Mike Markkula (the same man that had bankrolled the small company back in the mid 70s). IBM's new president Lou Gerstner headed Big Blue's team. They settled down to business. The IBM people gave a presentation on why they thought the deal would work, and how they would get the two companies' technology together.

Then it was the Apple people's turn to give their presentation. But before they did so, Gerstner called Spindler and Markkula into a side room with IBM's CFO and Jim Cannavino, head of IBM's PC division. They spoke behind closed doors for 45 minutes. When they emerged, Gerstner told his colleagues to pack up. The deal had fallen through, literally at the eleventh hour.

The sticking point was money. IBM offered Apple $US40 per share, a small premium over the stock price that valued Apple at around $US5 billion. Apple wanted $US7.5 billion, half as much again. IBM was prepared to compromise, Apple was not. Not for the first time, nor the last, Apple's greed and arrogance ensured its continued isolation from the rest of the industry.

If Apple and IBM had merged on that autumn day in 1994 the IT industry would today be a vastly different place. While the two would not have stopped Microsoft's rise to power – that was already well under way – they would have presented a viable alternative with the critical mass needed to compete in the corporate environment. As it is, IBM's ditched its OS/2 operating system shortly afterwards, and Apple retreated further into its specialist graphics and education markets.

The deal's failure received little publicity at the time. Both companies did an excellent job of keeping the negotiations quiet, and details only emerged subsequently. But had it proceeded it would have been the biggest IT story of the decade, and possibly ever. Apple would have ceased to be an independent company, but the industry would have had a major architectural alternative to what has become a virtual Wintel monopoly. The Big Blue Apple would have been the world's largest PC company.

It's All About Software

Without software computers are useless. The story of the computer industry is as much about programmers and the software they write as it is about the hardware. Microsoft's success demonstrates that fact.

The most influential modern programmer was also a woman – Grace Hopper, a scientist who joined the US naval reserve during the war and who rose to become an Admiral. She died in 1992. Her pioneering work in developing compilers (she invented the word) in the 1950s led directly to the development of FORTRAN, the world's first computer language.

Early computers were programmed physically, with switches and cables. Then came machine language (zeroes and ones), then assembly language, which consisted of abstruse instructions. FORTRAN, invented in 1954 by a young IBM researcher named John Backus, enabled people to program computers using simple English-like instructions and mathematical formulas. It led to a number of other languages, the most successful of which was COBOL (Common Business-Oriented Language).

Up until 1969 IBM included the software with the computer. Hardware was so expensive they could afford to give the software away. That changed in 1969, when IBM decided to "unbundle" software. Now others could compete against IBM, and the software industry was born. Independent software vendors (ISVs) proliferated, writing every type of application imaginable.

The first applications were accounting software. Everybody had a general ledger, after all. But most of these early systems had to be written from scratch for every different computer or operating system. That changed in 1970 when IBM's Ted Codd devised the concept of the relational database, which standardised the way computers stored and manipulated data.

A very different, but equally important, advance came in 1979 when Dan Bricklin invented Visicalc, the world's first spreadsheet program. Initially written for the Apple II, Visicalc and its many imitators revolutionised accounting and financial management to the extent that today it is impossible to imagine a world without spreadsheets.

The biggest battles in the history of computing have been over software standards. In 1984 Apple released the Macintosh, a revolutionary microcomputer that popularised the concept of the graphical user interface (GUI). GUIs used a mouse and the point-and-click concept, which made computers vastly easier to use. Eventually the whole industry followed suit, but not before massive legal battles over who had invented the idea.

The main battleground in the 1990s was over operating systems. Unix, invented at Bell Labs as a general purpose operating system in 1969, gradually supplanted older proprietary systems. But it was in turn challenged by Microsoft's Windows, a GUI-based system originally conceived for the desktop but increasingly capable of running on larger machines.

Windows won out over other desktop environments, totally destroying IBM's superior OS/2 and never allowing the Mac more than a small share of the market, but it has not been so successful at the high end. Its biggest challenge now comes from open source upstart Linux, an offshoot of Unix that offers an alternative both in technology and philosophy.

The GUI was invented at Xerox's Palo Alto Research Center (PARC) a smallish establishment on the northern fringes of Silicon Valley. It was established by the photocopier company in 1970 as a centre for pure research, its scientists free to develop whatever they wanted. And develop they did.

Other significant inventions to come out of PARC include Ethernet (which remains the standard computer networking technology), the computer mouse, the laser printer, object oriented programming, and the scientific workstation. Xerox never itself commercialised any of these products. PARC has acted as a kind of giant research lab for the rest of the industry, training a generation of scientists how to innovate and execute.

Networking the World

In the late 1960s the world's biggest computer user was the US Department of Defense. It had many machines of its own, and it used many more at universities and research institutions. Bob Taylor, a manager at ARPA (Advanced Research Projects Agency) proposed that ARPA's computers should be connected in some way, and he eventually persuaded ARPA to call for tenders.

A small company in Boston called BBN wrote a proposal for "interface message processors for the ARPA network". The company got the job, and BBN's Frank Hearst and his team started work, using a new technology called packet switching.

Packet switching sent data in discrete packets, rather than all at once. BBN was the first company to implement the technology, but the concept was also used picked up by a young man in Hearst's group called Bob Metcalfe, who four years later used it to devise Ethernet, the technology underlying most local area networks (LANs).

On 1 October 1969 the first Internet message was sent, from UCLA to Stanford Research Institute. Before the end of the year the University of California at Santa Barbara and the University of Utah were connected, in a network called ARPAnet.

Growth was slow. A year later the network had expanded to 15 nodes, but it took a further seven years to reach 100. Usage was restricted to academia and the military, and it remained very difficult to use until a networking standard called TCP/IP (Transmission Control Protocol/Internet Protocol) was developed by ARPA in 1982.

Gradually people began calling the new network the Internet. Its most widespread application became email, and things improved substantially in 1983 with the introduction of domain names, like .com and .org. But it was not until Tim Berners-Lee conceived the idea of the World Wide Web in 1989 that it began to resemble the Internet we know today.

Berners-Lee, an English scientist working at CERN, the European particle physics laboratory, came up with some simple specifications that made navigation around the Internet much easier. He devised a language called HTML (HyperText Markup Language), and a communications protocol called HTTP (HyperText Transfer Protocol) that used the concept of "hypertext" to allow people to jump easily between locations on the Internet.

But the Internet was still not a place for beginners. Addresses and locations were standardised, but you still had to know where you where going, and you needed a range of different software tools to get there. Enter the browser.

The browser was the brainchild of Marc Andreessen, a 21 year old student at the University of Illinois' National Center for Supercomputing Applications. Frustrated with how difficult it was to use the Internet, he enlisted the aid of a colleague, Eric Bina, to make it easier.

Over three months in the winter of 1992-93 Andreessen and Bina developed a piece of software that could navigate through hypertext links with the click of a mouse, that could display graphics as well as text, and which had an attractive and easy to use interface. They called it Mosaic.

The first version was ready in February 1993, and in April they released it for widespread use. By the end of the year there were over a million users. The next year Andreessen and Jim Clark started Netscape to commercialise a browser of the same name. Then Microsoft entered the browser market and started the acrimonious "browser wars" of the late 1990s, which led directly to Microsoft's current legal problems.

The browser was directly responsible for a vast increase in Internet usage. In January 1993 there were just 50 commercial web sites on the Internet (Congress had only authorised commercial usage the previous year). A year later, there were more than 10,000, and a year after that more than 100,000. Amazon.com, started by Jeff Bezos in 1994 became the archetype of a new e-business model. The long tech boom had begun. This, said many, was the New Economy, where the old rules did not apply any more.

Tim Berners-Lee, inventor of the World Wide Web, has not been idle in recent years. He is now director of the Word Wide Web Consortium (W3C), the non-profit coordinating body for Web development. His work still involves conceptualising where the Web is headed, and how to get it there.

Berners-Lee believes the next big thing will be the "Semantic Web", which he describes as an extension of the current Web where information is given meaning and where computers can not only process information but understand it.

The Web as it is currently constituted is optimised for human beings – to allow people easy access to documents. In the Semantic Web, data contained in Web pages will be coded with an extra dimension of information that will enable computers to make sense of it. We are part of the way there, with XML (Extensible Markup Language – an extension of HTML) and emerging Web Services protocols, but the Semantic Web will contain much more meaning. It will enable intelligent software agents to perform many of the searches and conduct many of the transactions that can currently only be undertaken by humans.

“Machines will become capable of analysing all the data on the Web – the content, links, and transactions between people and computers,” explains Burners-Lee. “The day-to-day mechanisms of trade, bureaucracy, and our daily lives will be handled by machines talking to machines, leaving humans to provide the inspiration and intuition.”

If this sounds far-fetched or difficult to imagine, just think how you would have explained today’s World Wide Web to someone twenty years ago.

The Information Millennium

On 10 March 2000 the NASDAQ index of US technology stock hit a peak of 5048.62, after doubling in the previous year. Today it is in the 1800s, a fall of 60 per cent from its peak. In July 2002 it was hovering in the low 1100s.

Since the top of the boom nearly 5000 Internet companies in the US, and many more internationally, have gone broke or been acquired. 200,000 jobs were lost in Silicon Valley alone. The market capitalisation of all the world’s computer companies is about 20 per cent of what it was three and half years ago.

After every boom there is a bust. Rarely has this fact been more graphically demonstrated than in the Great Tech Bust of recent years. While some optimistic souls are waiting for things to return to “normal”, most people have come to the realisation that what we now have is normal, and that it was the tech boom that was the aberration.

The last ten years have seen the information technology industry change beyond recognition. Large and proud companies like DEC, Wang, Compaq, Prime, Data General, Amdahl and many others are no more. PCs have become so cheap that they are commodities, and so powerful that no-one takes much notice of their technical specifications.

The Internet is a fact of life. Most homes are connected to the Net, and it is impossible to do business without it. Most airline tickets are bought over the Internet, email is the standard method of commercial communication, and all media are digital. Mobile phones and other wireless devices are commonplace, removing physical location as a constraint on communication. There is more processing power in the average car than in the largest data centre of the 1970s, and even our fridges are online (some of them, anyway).

Yet, all around us, there is talk of doom and gloom. Oracle’s Larry Ellison famously said recently that 90 per cent of the companies in Silicon Valley don’t deserve to exist, and even Gartner says that half them will go out of business in the next five years (0.7 probability factor). We are awash with technology, but the industry is crying poor.

The reason is very simple. The IT industry, after more than 50 years of fabulous growth, has grown up. It has matured. As its products have commoditised its margins have reduced. The offerings from one player look much like those from another. In most parts of the industry the number of significant companies has fallen to just three or four.

One company, Microsoft, has a virtual monopoly on desktop software – operating systems and applications. Its abuse of its position has led to legal action by the US Department of Justice, which went after IBM in the 1970s for just the same reason. The IBM case eventually fell apart, not because IBM had not abused its position (it had), but because the industry had changed so much that the circumstances had become irrelevant.

The same is happening in Microsoft's case. The software giant has been exposed as a serial bully and a master of IBM's old FUD (fear, uncertainty and doubt) tactics, but it is all becoming meaningless. Microsoft's dominance will decline, because of the very market forces it claims to uphold.

The ubiquity of the Internet will ensure that. Software, and the functions it performs, are increasingly traded as Web-delivered services. The open source movement, which is already making significant inroads with Linux, is another force of the future. The very concept of intellectual property, be it software, or music, or film, is under threat from the new technology. That will be the battleground of the future.

We are moving towards the era of the information utility, with IT delivered on demand via an invisible grid that encircles the globe. If the last half century has been fun, the next will be fascinating.

Email	**graeme@philipson.info**
Web	**www.philipson.info**
Ph	**+61 2 4226 2200**
Fax	**+61 2 4226 2201**
Mobile	**+61 418 609 397**
Post	**PO Box 290, Figtree NSW 2525**
Street	**147 Koloona Ave, Mt Keira NSW 2500** **AUSTRALIA**