



The  
Economist

# THE FUTURE OF TECHNOLOGY



THE ECONOMIST IN ASSOCIATION WITH  
PROFILE BOOKS LTD

Published by Profile Books Ltd  
3A Exmouth House, Pine Street, London EC1R 0JH

Copyright © The Economist Newspaper Ltd 2005

All rights reserved. Without limiting the rights under copyright reserved above, no part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form or by any means (electronic, mechanical, photocopying, recording or otherwise), without the prior written permission of both the copyright owner and the publisher of this book.

The greatest care has been taken in compiling this book.  
However, no responsibility can be accepted by the publishers or compilers  
for the accuracy of the information presented.

Where opinion is expressed it is that of the author and does not necessarily coincide  
with the editorial views of The Economist Newspaper.

Typeset in EcoType by MacGuru  
info@macguru.org.uk



Printed and bound in Great Britain by  
Creative Print and Design (Wales), Ebbw Vale

A CIP catalogue record for this book is available  
from the British Library

ISBN 1 86197 971 1

# Contents

The authors	vi
Foreword	vii
<b>Part 1 Information technology grows up</b>	
1 Coming of age	3
2 Securing the cloud	41
3 Make it simple	77
4 A world of work	111
<b>Part 2 The shift towards consumer electronics</b>	
5 Mobile phones	149
6 Gaming	185
7 Digital home	201
<b>Part 3 Searching for the next big thing</b>	
8 Climbing the helical staircase	235
9 Energy	273
10 Small wonders	305
11 Robots and artificial intelligence	331
Index	341



# Foreword

**T**O UNDERSTAND THE FUTURE OF TECHNOLOGY, start by looking at its past. From the industrial revolution to the railway age, through the era of electrification, the advent of mass production and finally to the information age, the same pattern keeps repeating itself. An exciting, vibrant phase of innovation and financial speculation is followed by a crash; then begins a longer, more stately period during which the technology is widely deployed. Consider the railway mania of the 19th century, the dotcom technology of its day. Despite the boom and bust, railways subsequently proved to be a genuinely important technology and are still in use today – though they are not any longer regarded as technology by most people, but as merely a normal part of daily life.

Having just emerged from its own boom and bust, the information-technology industry is at the beginning of its long deployment phase. Irving Wladawsky-Berger, a technology guru at IBM, has said that the industry has entered its “post-technology” period. His point is not that technology no longer matters, but that how it is applied is now more important than the technology itself. All successful technologies eventually fade into the background of everyday life, as electricity and railways have, becoming both more important and less visible. That is now happening to information technology. The history of technology thus provides clues about its future.

To be sure, the information-technology industry is ambivalent about its growing maturity. On the one hand, it means that the debates of the past decade – whether businesses have any choice but to embrace the internet, for example, or whether it is possible to make money online – have been resolved, and in the technologists’ favour. All the hype of the internet boom contained a kernel of truth, though harnessing the new technology proved harder and is taking longer than the cheerleaders of the 1990s anticipated. That is the nature of revolutions: they are rarely as quick or as clean as the revolutionaries expect them to be.

On the other hand, however, this means that the revolutionary ideas of a few years ago have now become conventional wisdom. Having convinced the business world of the merits of technology, the industry has lost much of its iconoclastic fervour. The corporate adoption of information technology has become such old hat, indeed, that Nicholas Carr, an editor at the *Harvard Business Review*, even published an article

in 2003 entitled “IT doesn’t matter”. Many technologists were livid, but Mr Carr had a point. While there are some areas where technology can still provide competitive advantage, its widespread adoption means that it makes less difference than it used to.

This shift has many implications for the future of technology, which will be examined in the first part of this book. For one thing, the industry needs to adjust its self-image as it moves from youthful exuberance to middle-aged pragmatism. It also needs to place more emphasis on security and reliability, as computers become central to so many aspects of business and personal life, from e-mail to online shopping and banking. Fixing security problems may be less sexy than dreaming up whizzy new technologies, but it must be done.

Another problem is that while information technology can do all sorts of clever things, it is still too complex in many respects, and places insufficient emphasis on the needs of its users. This is understandable in the early days, when the main thing is to get new technologies to work at all, but eventually it hinders their spread and limits their potential. Lastly, companies are increasingly exploiting the plunging costs of computing and communications to make their internal processes more flexible and efficient, in particular through outsourcing. This has put information technology at the heart of the battles over globalisation and development – unfamiliar ground for most technologists, who generally prefer to avoid politics.

Does this mean that technology has completely lost its shine and become a dull, tedious business? Far from it. Rather, as corporate computing has matured, the technology industry’s innovative focus has moved into consumer electronics. It used to be the case that the most advanced computing and communications technologies were to be found sealed away in corporate data-centres; but new technologies now appear first in consumer gadgets, which are produced by the million. Mobile phones, nearly 2 billion of which are now in use around the world, have become the most ubiquitous form of information technology on earth, while video-game consoles have become one of the most advanced, with far more computing power than office PCs.

The growing ubiquity and sophistication of consumer-electronics devices is the topic of the second part of this book. Consumer-electronics firms are taking the fruits of the information-technology boom and applying them not to corporate systems, where they started out, but to every conceivable aspect of daily life. It is a logical consequence of the information-technology boom of the late 20th century. The first hard

disks, for example, were attached to mainframes; now hard disks can be found in the living room, inside personal video recorders, or even in your pocket, inside an iPod music-player.

As well as mobile phones and games consoles, the two areas where the rapid pace of innovation is most evident, there is an industry-wide effort to interconnect consumer-electronics devices of all kinds in a seamless way, to create a “digital lifestyle” within a “digital home”. It is still more of a vision than a reality, but the building blocks from which the digital home will be constructed are becoming apparent: exotic new wireless technologies, cheap storage, broadband internet access and large, flat-panel displays. All of this ought to allow music, photos, movies and games to be accessed and enjoyed with unprecedented ease – eventually. But first there are technical hurdles to overcome, new business models to devise, and legal and copyright problems to sort out. For the music and movie industries are institutionally suspicious of new technologies – despite the fact that, time after time, new inventions that they initially perceived as threats ended up expanding their businesses.



In a sense, though, the current boom in consumer technology is merely an aftershock of the information-technology boom, rather than an entirely new phase of technological development. So what comes next? Once information technology has percolated into everything, from wireless sunglasses that double as headphones to radio-tagged cereal boxes, what new technology will lead the next great phase of transformation, disruption and creativity?

The third part of this book will examine some of the leading contenders. The first is biotechnology, which promises new medical treatments tailored for individual patients, new crops and novel industrial processes. Indeed, the industrial uses of genetic modification could prove be far more significant than its better known (and highly controversial) agricultural applications. A second contender is energy technology. There have been energy revolutions in the past, notably the introduction of steam power, which spawned the industrial revolution, and the advent of electrification, which prompted further industrial transformation. As global energy consumption grows, reserves of fossil fuels decline and climate change accelerates, there is likely to be a growing need for better, cleaner energy technologies.

A third candidate is nanotechnology, the exploitation of the unusual phenomena that manifest themselves at the nanoscale (a nanometre is a thousand millionth of a metre). This emerging field has been the topic of speculation and controversy for some years, but is now moving from

the realm of science fiction to the marketplace. Lastly, we consider two technologies - robotics and artificial intelligence - that have been touted as the “next big thing” in the past and are now considered failures. Nevertheless, both have arguably been quite widely adopted, which suggests that evaluating the success of new technologies is more difficult than it seems.

All three parts of this book consist of surveys and articles that appeared in *The Economist* between 2002 and 2005; the date of publication is given at the end of each one. Each article reflects the author’s point of view at the time in question, but while some have been lightly edited, they have not been substantially rewritten, and are as valid today as they were at the time of publication. Collectively they illustrate how the technology industry is changing, how technology continues to affect many aspects of everyday life - and, looking further ahead, how researchers in several promising fields are developing the innovations that seem most likely to constitute the future of technology.



## Paradise lost

**So far, information technology has thrived on exponentials. Now it has to get back to earth**

CLOSE YOUR EYES and think of information technology. You might picture your PC crashing yet again, and recall that your teenager was supposed to fix it. That leads you to the 12-year-old hacker who broke into a bank's computer system the other day, which brings to mind the whizz-kids in a garage inventing the next big thing that will turn them into the youngest billionaires ever.

In IT, youth seems to spring eternal. But think again: the real star of the high-tech industry is in fact a grey-haired septuagenarian. Back in 1965, Gordon Moore, co-founder of Intel, the world's biggest chipmaker, came up with probably the most famous prediction in IT: that the number of transistors which could be put on a single computer chip would double every 18 months. (What Mr Moore actually predicted was that the figure would double every year, later correcting his forecast to every two years, the average of which has come to be stated as his "law".)

This forecast, which implies a similar increase in processing power and reduction in price, has proved broadly accurate: between 1971 and 2001, transistor density has doubled every 1.96 years (see Chart 1.1). Yet this pace of development is not dictated by any law of physics. Instead, it has turned out to be the industry's natural rhythm, and has become a self-fulfilling prophecy of sorts. IT firms and their customers wanted the prediction to come true and were willing to put up the money to make it happen.

Even more importantly, Moore's law provided the IT industry with a solid foundation for its optimism. In high-tech, the mantra goes, everything grows exponentially. This sort of thinking reached its peak during the internet boom of the late 1990s. Suddenly, everything seemed to be doubling in ever-shorter time periods: eyeballs, share prices, venture capital, bandwidth, network connections. The internet mania began to look like a global religious movement. Ubiquitous cyber-gurus, framed by colourful PowerPoint presentations reminiscent of stained glass, prophesied a digital land in which growth would be limitless, commerce frictionless and democracy direct. Sceptics were derided as bozos "who just don't get it".



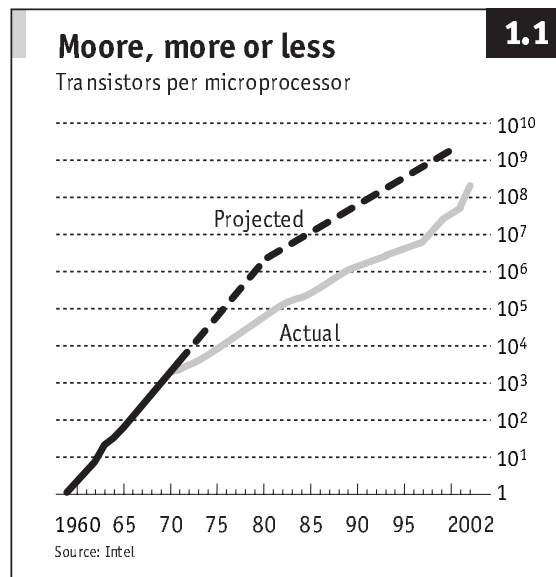
Today, everybody is older and wiser. Given the post-boom collapse of spending on IT, the idea of a parallel digital universe where the laws of economic gravity do not apply has been quietly abandoned. What has yet to sink in is that this downturn is something more than the bottom of another cycle in the technology industry. Rather, the sector is going through deep structural changes which suggest that it is growing up or

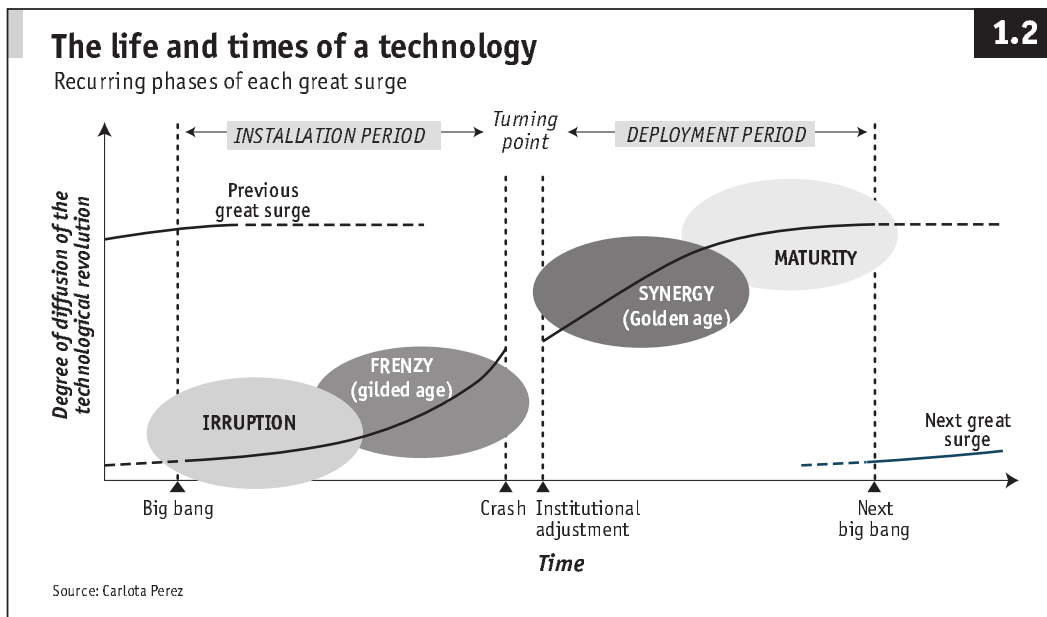
even, horrors, maturing. Silicon Valley, in particular, has not yet come to grips with the realities, argues Larry Ellison, the chief executive of Oracle, a database giant (who in his early 60s still sports a youthful hairdo). “There’s a bizarre belief that we’ll be young forever,” he says.

It is not that Moore’s law has suddenly ceased to apply. In fact, Mr Moore makes a good case that Intel can continue to double transistor density every 18 months for another decade. The real issue is whether this still matters. “The industry has entered its post-technological period, in which it is no longer technology itself that is central, but the value it provides to business and consumers,” says Irving Wladawsky-Berger, a senior manager at IBM and another grey-haired industry elder.

Scholars of economic history are not surprised. Whether steam or railways, electricity or steel, mass production or cars – all technological revolutions have gone through similar long-term cycles and have eventually come of age, argues Carlota Perez, a researcher at Britain’s University of Sussex, in her book *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages* (Edward Elgar, 2002).

In her model (see Chart 1.2 overleaf), technological revolutions have two consecutive lives. The first, which she calls the “installation period”, is one of exploration and exuberance. Engineers, entrepreneurs and investors all try to find the best opportunities created by a technological big bang, such as Ford’s Model T in 1908 and Intel’s first microprocessor in 1971. Spectacular financial successes attract more and more capital, which leads to a bubble. This is the “gilded age” of any given technology, “a great surge of development”, as Ms Perez calls technological revolutions.





The second, or “deployment”, period is a much more boring affair. All the quick bucks have been made, so investors prefer to put their money into the real economy. The leading firms of the new economy become bigger and slower. The emphasis is no longer on raw technology, but on how to make it easy to use, reliable and secure. Yet this period is also the “golden age” of a technology, which now penetrates all parts of society.

These two periods of a technological revolution are separated by what Ms Perez calls a “turning point” – a crucial time for making the choices that determine whether a technological revolution will deliver on its promises. In her book, she concentrates mainly on the social and regulatory decisions needed to allow widespread deployment of new technology. But the same argument applies to technology vendors and customers. To enter their “golden age”, they have to leave their youthful excesses behind and grow up.

### A duller shade of gold

This section will examine how much grey the IT industry (and its leaders’ hair) has already acquired. The first three articles are about technological shifts, and how value is moving from the technology itself to how it is applied. Many of the wares that made the IT industry’s fortunes in the installation period are becoming a commodity. To overcome this problem, hardware vendors are developing new software that allows networks of machines to act as one, in effect turning com-

puting into a utility. But the IT industry's most profitable layer will be services of all kinds, such as software delivered as an online service, or even business consulting.

The second half of this section looks at institutional learning, which has caused the value created by the IT industry to be increasingly captured by its customers. For the first time in its history, the IT industry is widely adopting open standards. Equally important, buyers are starting to spend their IT budgets more wisely. Meanwhile, the industry's relationship with government is becoming closer.

All this suggests that the technology industry has already gone greyish at the temples since the bubble popped, and is likely to turn greyer still. Sooner or later the sector will enter its "golden age", just as the railways did. When Britain's railway mania collapsed in 1847, railroad shares plunged by 85%, and hundreds of businesses went belly-up. But train traffic in Britain levelled off only briefly, and in the following two decades grew by 400%.

So are the IT industry's best days yet to come? There are still plenty of opportunities, but if the example of the railways is anything to go by, most IT firms will have to make do with a smaller piece of the pie. As this newspaper (then called *The Economist*, *Weekly Commercial Times*, *Bankers' Gazette*, and *Railway Monitor*) observed in 1857: "It is a very sad thing unquestionably that railways, which mechanically have succeeded beyond anticipation and are quite wonderful for their general utility and convenience, should have failed commercially."

Brad DeLong, an economics professor at the University of California at Berkeley, puts it somewhat more succinctly: "I am optimistic about technology, but not about profits."

## Modifying Moore's law

**Many of the innovations that made the IT industry's fortunes are rapidly becoming commodities - including the mighty transistor**

IF GOOGLE WERE TO CLOSE down its popular web-search service tomorrow, it would be much missed. Chinese citizens would have a harder time getting around the Great Firewall. Potential lovers could no longer do a quick background check on their next date. And college professors would need a new tool to find out whether a student had quietly lifted a paper from the internet.

Yet many IT firms would not be too unhappy if Google were to disappear. They certainly dislike the company's message to the world: you do not need the latest and greatest in technology to offer outstanding services. In the words of Marc Andreessen of Netscape fame, now chief executive of Opsware, a software start-up: "Except applications and services, everything and anything in computing will soon become a commodity."

Exactly what is meant by "commoditisation", though, depends on whom you talk to. It is most commonly applied to the PC industry. Although desktops and laptops are not a truly interchangeable commodity such as crude oil, the logo on a machine has not really mattered for years now. The sector's most successful company, Dell, is not known for its technological innovations, but for the efficiency of its supply chain.

As the term implies, "commoditisation" is not a state, but a dynamic. New hardware or software usually begins life at the top of the IT heap, or "stack" in geek speak, where it can generate good profits. As the technology becomes more widespread, better understood and standardised, its value falls. Eventually it joins the sector's "sediment", the realm of bottom feeders with hyper-efficient metabolisms that compete mainly on cost.

### **Built-in obsolescence**

Such sedimentation is not unique to information technology. Air conditioning and automatic transmission, once selling points for a luxury car, are now commodity features. But in IT the downward movement is much faster than elsewhere, and is accelerating - mainly thanks to

Moore's law and currently to the lack of a new killer application. "The industry is simply too efficient," says Eric Schmidt, Google's chief executive (who seems to have gone quite grey during his mixed performance at his previous job as boss of Novell, a software firm).

The IT industry also differs from other technology sectors in that its wares become less valuable as they get better, and go from "undershoot" to "overshoot," to use the terms coined by Clayton Christensen, a professor at Harvard Business School. A technology is in "undershoot" when it is not good enough for most customers, so they are willing to pay a lot for something that is a bit better although not perfect. Conversely, "overshoot" means that a technology is more than sufficient for most uses, and margins sink lower.


PCs quickly became a commodity, mainly because IBM outsourced the components for its first venture into this market in the early 1980s, allowing others to clone the machines. Servers have proved more resistant, partly because these powerful data-serving computers are complicated beasts, partly because the internet boom created additional demand for high-end computers running the Unix operating system.

But although expensive Unix systems, the strength of Sun Microsystems, are - and will probably remain for some time - a must for "mission-critical" applications, servers are quickly commoditising. With IT budgets now tight, firms are increasingly buying computers based on PC technology. "Why pay \$300,000 for a Unix server," asks Mr Andreessen, "if you can get ten Dell machines for \$3,000 each - and better performance?"

Google goes even further. A visit to one of the company's data centres in Silicon Valley is a trip back to the future. In the same way that members of the Valley's legendary Homebrew Computer Club put together the first PCs using off-the-shelf parts in the early 1970s, Google has built a huge computer system out of electronic commodity parts.

### **Modern Heath Robinsons**

When the two Stanford drop-outs who founded Google, Sergey Brin and Larry Page, launched the company in 1998, they went to Fry's, an electronics outlet where the Valley's hardcore computer hobbyists have always bought their gear. Even today, some of the data centres' servers appear to be the work of tinkerers: circuit boards are sagging under the weight of processors and hard disks, and components are attached by Velcro straps. One reason for the unusual design is that parts can be easily swapped when they break. But it also allows Google's servers to



be made more powerful without having to be replaced completely.

What makes it easier for Google to swap off-the-shelf components is that much of its software is also a commodity of sorts. Its servers run Linux, the increasingly popular open-source operating system developed by a global community of volunteer programmers, and Apache, another open-source program, which dishes up

web pages.

Because Google has always used commodity hardware and software, it is not easy to calculate how much money it has saved. But other firms that have recently switched from proprietary gear say they have significantly reduced their IT bill. Amazon.com, the leading online shopping mall, for instance, managed to cut its quarterly technology spending by almost \$20m (see Chart 1.3).

The most interesting feature of Google's data centre, however, is that its servers are not powered by high-end chips, and probably will not have Itanium, Intel's most powerful processor, inside for some time yet, if ever. This sets Google apart among hot Silicon Valley start-ups, whose business plans are mostly based on taking full advantage of the exponential increase in computing power and similar growth in demand for technology.

"Forget Moore's law," blared the headline of an article about Google in *Red Herring*, a now-defunct technology magazine. That is surely overblown, but Google's decision to give Itanium a miss for now suggests that microprocessors themselves are increasingly in "overshoot", even for servers - and that the industry's 30-year race for ever more powerful chips with smaller and smaller transistors is coming to an end.

Instead, other "laws" of the semiconductor sector are becoming more important, and likely to change its underlying economics. One is the fact that the cost of shrinking transistors also follows an exponential upward curve. This was no problem as long as the IT industry gobbled up new chips, thus helping to spread the cost, says Nick Tredennick,

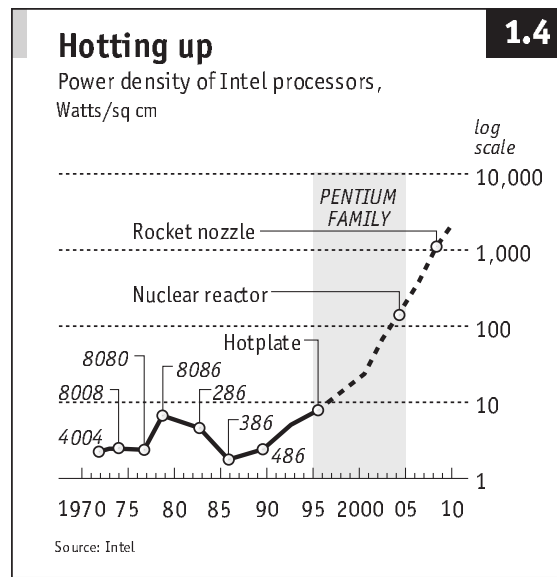
editor of the *Gilder Technology Report*, a newsletter. But now, argues Mr Tredennick, much of the demand can be satisfied with “value transistors” that offer adequate performance for an application at the lowest possible cost, in the same way as Google’s. “The industry has been focused on Moore’s law because the transistor wasn’t good enough,” he says. “In the future, what engineers do with transistors will be more important than how small they are.”

This is nothing new, counters Paul Otellini, Intel’s chief executive. As chips become good enough for certain applications, new applications pop up that demand more and more computing power, he says: once Google starts offering video searches, for instance, it will have to go for bigger machines. But in recent years, Intel itself has shifted its emphasis somewhat from making ever more powerful chips to adding new features, in effect turning its processors into platforms.

In 2003 it launched Centrino, a group of chips that includes wireless technology. The Centrino chips are also trying to deal with another, lesser-known, limiting factor in chipmaking: the smaller the processors become, the more power-hungry and the hotter they get (see Chart 1.4). This is because of a phenomenon called leakage, in which current escapes from the circuitry. The resulting heat may be a mere inconvenience for users of high-end laptops, who risk burning their hands or thighs, but it is a serious drawback for untethered devices, where it shortens battery life – and increasingly for data centres as well, as Google again shows.



### Cool chips

The firm’s servers are densely packed to save space and to allow them to communicate rapidly. The latest design is an eight-foot rack stuffed with 80 machines, four on each level. To keep this computing powerhouse from overheating, it is topped by a ventilation unit which sucks air through a shaft in its centre. In a way, Google is doing to servers



what Intel has done to transistors: packing them ever more densely. It is not the machines' innards that count, but how they are put together.

Google has thus created a new computing platform, a feat that others are now replicating in a more generalised form. Geoffrey Moore (no relation), chairman of the Chasm Group, a consultancy, and a partner at Mohr, Davidow Ventures, a Silicon Valley venture-capital firm, explains it this way: computing is like a game of Tetris, the computer-game classic; once all the pieces have fallen into place and all the hard problems are solved, a new playing field emerges for others to build on.





## Moving up the stack

### The network is becoming the computer - and the IT industry's dominant platform

COMPUTING IS SUPPOSED TO BE the ultimate form of automation, but today's data centres can be surprisingly busy with people. When an application has to be updated or a website gets more visitors than expected, system administrators often have to install new programs or set up new servers by hand. This can take weeks and often turns out to be more complicated than expected.

Google's data centres, however, look deserted most of the time, with only about 30 employees to look after a total of 54,000 servers, according to some estimates. This is in part because machines doing searches need less care than those running complex corporate applications; but more importantly, the firm's programmers have written code that automates much of what system administrators do. It can quickly change a computer that sifts through web pages into a server that dishes up search results. Without the program, Google would have to hire many more people.

It all goes to show that another law in computing, proclaimed by Gordon Bell, another greying industry legend, still holds true: in IT, the dominant platform shifts every ten years or so. Mainframes, minicomputers, PCs and servers are now likely to be followed by a grid of computers, either within a data centre or as a disparate collection of connected machines. The network will at last be the computer, to paraphrase a slogan coined by Sun Microsystems. Machines will no longer simply be attached to a network: instead, the network will allow them to act as one.

Yet this new platform, which computer scientists like to call "grid computing", is less about replacing old technology and more about managing the existing gear - another sign that IT is maturing. Merrill Lynch's Steve Milunovich, one of the leading hardware analysts on Wall Street, says that IT has entered the era of "managed computing". Forrester Research, a high-tech consultancy, has coined the term "organic IT" - a computing infrastructure that is not only built on cheap parts, but is also as adaptive as a living organism. Whatever label the industry settles for, the race to lead in the next round of computing is already on. The new

platform gives those threatened by commoditisation a chance to differentiate themselves by moving up the technology stack to a potentially more lucrative layer.

There is every incentive for HP, IBM, Microsoft and Sun, as well as a raft of start-ups, to encourage this shift, but there is also a real need for a new platform. Computing has certainly got faster, smarter and cheaper, but it has also become much more complex. Ever since the orderly days of the mainframe, which allowed tight control of IT, computer systems have become ever more distributed, more heterogeneous and harder to manage.

### **Managing complexity**

In the late 1980s, PCs and other new technologies such as local area networks (LANs) allowed business units to build their own systems, so centralised IT departments lost control. In the late 1990s, the internet and the emergence of e-commerce “broke IT’s back”, according to Forrester. Integrating incompatible systems, in particular, has become a big headache.

A measure of this increasing complexity is the rapid growth in the IT services industry. According to some estimates, within a decade 200M IT workers will be needed to support a billion people and businesses connected via the internet. Managing a storage system already costs five times as much as buying the system itself, whereas less than 20 years ago the cost of managing the system amounted to only one-third of the total (see Chart 1.5).

What is more, many of today’s IT systems are a patchwork that is inherently inefficient, so firms spend 70–90% of their IT budgets simply on keeping their systems running. And because those systems cannot adapt quickly to changes in demand, companies overprovide. They now spend almost \$50 billion a year on servers, but the utilisation rate for these computers is often below 30%.

Besides, complexity is bound to increase, predicts Greg Papadopoulos, Sun’s chief technology officer. Today, the electronics to hook up any device to the network cost about \$1. In ten years’ time, the price will be down to one cent. As a result, he says, the number of connected things will explode, and so will the possible applications. For example, it will become practical to track items such as razor blades (10% of which apparently disappear on their way from the factory to the shop).

When things get too complicated, engineers usually add a layer of code to conceal the chaos. In some ways, the current shift in computing

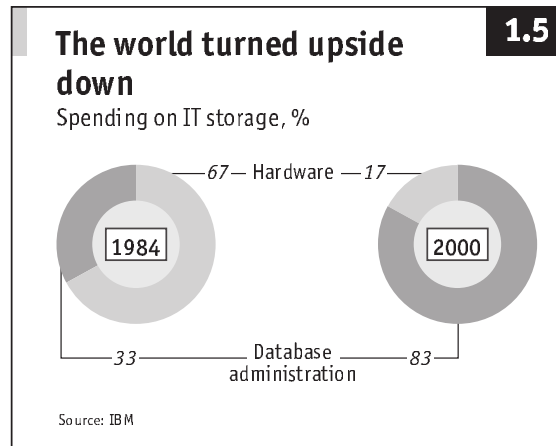
is the equivalent of what happened when cars became easy to use and drivers only had to turn the key instead of having to hand-crank the engines. In geek speak, adding such a new layer is called “raising the level of abstraction”. This happened when PC operating systems first hid the nuts and bolts of these computers and gave them a simple user interface, and it is happening now with the new platform, which is already

being compared to an operating system for data centres or computing grids.

Just like Google’s management program, this grid computing software (only half-jokingly called “griddleware” by some) automates much of the work of system administrators. But it is also supposed to serve a higher purpose: “virtualisation”. Put simply, this means creating pools of processing power, storage capacity and network bandwidth. A data centre, or a collection of machines on a network, thus becomes a virtual computer whose resources can be allocated as needed. The ultimate goal is that managed computing will become rather like flying a modern jet plane: IT workers will tell the system what kind of applications it should run, and then deal only with exceptions.



Although the rivals in this new field are pretty much on the same technological track, their strategies are different. Some of the numerous start-ups already have working products – and no hidden agenda, says Mr Andreesen, of Opsware, the leading newcomer: “We don’t need to push our customers also to buy other stuff from us.” The incumbents, on the other hand, want the new software layer to protect their old business models as well. HP’s Utility Data Centre (UDC) initiative and Sun’s N1 plan are supposed to help these firms sell their profitable hardware. IBM’s “autonomic computing” effort goes hand-in-hand with Big Blue’s IT services business. And Microsoft’s Dynamic Services Initiative (DSI) is tightly linked with its Windows operating system.

Yet despite such arm-twisting, customers are unlikely to bet solely on newcomers. Only the biggest vendors will really be able to deliver managed computing, argues Shane Robinson, the chief technology officer of HP, which has much riding on the new platform. According to the



Gartner Group, a consultancy, HP is leading in virtualisation, and views management software as its big opportunity.

Once thing is clear: once all the technical challenges of grid computing have been overcome, hardware will have become a true commodity. Machines, storage devices and networks will lose their identity and feed into pools of resources that can be tapped as needed. This liquefaction of hardware, in turn, will allow computing to become a utility, and software a service delivered online.



# Techniques, not technology

## IT firms hope to turn the dismal science into a profitable business

**T**O MOST PEOPLE the world is round, but geeks often see it as a stack of layers. In corporate computing, it starts with the hardware, on top of which sits the operating system, then the database, the applications and finally IT services. When their layer is getting commoditised, technology companies tend to move up the stack, where more money can be made.

In their quest for greener pastures, IT firms have reached new heights by moving into cutting-edge economics. Both HP and IBM have opened labs to conduct research on the subject, in the hope that this will help them to offer their customers more sophisticated services.

To be sure, economics has had its place in the IT industry for some years now. HP, for instance, already uses software that simulates markets to optimise the air-conditioning systems in its utility data centres. And IBM's Institute for Advanced Commerce has studied the behaviour of bidding agents, in the hope of designing them in such a way that they do not engage in endless price wars.

Now HP is reaching even higher, with experimental economics. As the name implies, researchers in this field set up controlled experiments with real people and real money to see whether economic theories actually work. Perhaps surprisingly, it seems that they do, as demonstrated by the work of Vernon Smith of George Mason University in Virginia. (Mr Smith is considered the founding father of this field and won the 2002 Nobel prize in economics.)

### Secret agent



HP goes further. The firm's team of five researchers does not test economic theories, but tries to create "novel mechanisms to improve the fluidity of interactions in the information economy", says Bernardo Huberman, head of the group. In everyday language, the researchers are working on clever tools that make it easier to negotiate online, establish reputations and make forecasts.

Mr Huberman's group already has something to show for its efforts. It has developed a methodology for predicting uncertain events using a small group of individuals. First, they find out about their subjects'

attitudes towards risk and their ability to forecast a given outcome. They then use this information to weight and aggregate their predictions of events, resulting in fairly accurate forecasts.

These tools will first be used inside the company. The top management of one of HP's divisions is already testing the forecasting methodology to predict its revenue. But ultimately the firm wants to find outside customers for its research findings. American intelligence agencies, such as the CIA, have already shown interest. They need better tools to weigh the opinions of those who analyse incoming information.

So at what point will firms such as HP and IBM have moved far enough up the stack to cease to be traditional IT vendors and become service providers or consultancies? Most analysts agree that this metamorphosis is still some way off. But it already seems certain that in future IT firms will increasingly be in the business of techniques rather than technology.



## At your service

### Despite early failures, computing will eventually become a utility

**M**ARC BENIOFF (not a grey hair in sight) is not afraid to mix religion and business. In February 2003, the practising Buddhist and chief executive of salesforce.com, a San Francisco start-up, invited 200 customers and friends to a benefit concert featuring David Bowie, with proceeds going to the Tibet House, a New York cultural centre whose patron is the Dalai Lama. But Mr Benioff also used the event to get his firm's message across: "Freedom from software".

The unusual mixture raised some eyebrows, but in a way Mr Benioff's technology does indeed dispense with software. Clients can access the firm's service via a web browser, which saves them having to install a complicated customer relationship management (CRM) program on their own computers. Some 6,300 customers in 110 countries have already signed up for it, generating \$52m in the 2002/03 fiscal year, says Mr Benioff.

Sceptics maintain that salesforce.com is not so much the leader of a new trend as a lone survivor of better times. Hundreds of application service providers (ASPs) were launched late in the dotcom boom, but few others have succeeded. This is mainly because with existing technology it is difficult to make money on a high-end software service.

But not for much longer. Thanks to the technological trends outlined on pages 8-16, computing is becoming a utility and software a service. This will profoundly change the economics of the IT industry. "The internet spells the death of the traditional software business model," predicts Mr Benioff.

This is not as earth-shattering as it sounds. As other technologies matured, buyers were given more choice in how to acquire them, says IBM's Irving Wladawsky-Berger. In the early days of electricity, for instance, most firms had to have their own generators. Now most can get their power from the grid. Similarly, he says, it would be surprising if in 20 years' time most of IT was not outsourced.

Traditionally, companies wanting to invest in computer systems did not have much choice: they had to build and operate their own. To be sure, they could outsource the work to companies such as EDS and IBM Global Services, but in technical terms that did not change much,

because such firms usually operate dedicated computer systems for each customer.

### **There must be a better way**

When it comes to enterprise software, in particular, this way of delivering technology creates a somewhat perverse set of economics. Software is a service at heart, albeit an automated one, but it is sold much like a manufactured good. Customers have to pay large sums of money up front, bear much of the risk that a program may not work as promised and cannot readily switch vendors.

IT firms, for their part, have to spend a lot of resources on marketing and distribution, rather than concentrating on developing software that works well and is easy to use. Network effects and Wall Street make matters worse. In many markets it is a great advantage to be first, so vendors are tempted to release programs even if they are still riddled with bugs. And because equity analysts rightly consider software firms a risky investment, such firms must grow quickly to justify their relatively high share prices, pushing them to sell more programs than customers need.

All this explains several peculiarities of the software business. One is the fact that many of the licences sold are never used, a phenomenon known as “shelfware”. More importantly, many software firms have grown so fast, often mortgaging the future, that they collapse when they reach \$1 billion in annual revenues, sometimes never to recover. Then there is the end-of-quarter rush, spurring many firms to do anything to get deals signed and meet analysts’ expectations.

The need to grow quickly also explains why IT industry leaders are such a “truly extraordinary cast,” in the words of Louis Gerstner, who was IBM’s chief executive for eight years. “They make outrageous remarks, they attack one another publicly with great relish,” he writes in his book *Who Says Elephants Can’t Dance?* (HarperBusiness 2002). Bosses of software firms, in particular, need to demonstrate that they will pursue growth at almost any price – which explains why they are often paired with grey-haired chief financial officers as a calming counterweight.

Mr Gerstner, who has spent most of his career outside the IT industry, does not point fingers, but examples of the industry’s “bizarre practices”, as he puts it, are not hard to find. The most obvious one is Oracle, a database giant which had a near-death experience in 1991, having cut some reckless deals to meet expectations. The firm is also known to



have released software prematurely, most recently its e-business suite. And it is run by Larry Ellison, arguably the most colourful boss in the industry, and Jeff Henley, the archetype of a grandfatherly CFO. To be fair, it must be said that the company has matured greatly in recent years.

In future, technology itself could lead to a better balance in the sector as a whole. The internet made it possible to run ASPs such as salesforce.com, but it also enabled hardware-makers to monitor servers and bill customers remotely on the basis of the average use per month. This is the sort of thing HP does with its high-end Superdome machines.

As data centres become automated, computing will increasingly turn into a true utility. With the management software described on pages 13-16, firms can share computing resources, which means they always have enough of them but pay only for what they actually use. They no longer need to run their own dedicated machines, any more than they need to run their own power plants.

### **Waiting for web services**

For software truly to become a service, however, something else has to happen: there has to be wide deployment of web services. These are not, as the term might suggest, web-based offerings such as salesforce.com, but a standard way for software applications to work together over the internet. Google, for instance, also offers its search engine as a web service to be used in other web offerings, such as Googlefight, a site where surfers with time to waste can find out which of two related key words produces more search results.

Ultimately, experts predict, applications will no longer be a big chunk of software that runs on a computer but a combination of web services; and the platform for which developers write their programs will no longer be the operating system, but application servers. These are essentially pieces of software that offer all the ingredients necessary to cook up and deliver a web service or a web-based service such as salesforce.com.

Just as with management software for data centres, vendors are already engaged in a battle for dominance. Ranged on one side is Microsoft with its .NET platform (although it has recently toned down the use of this name). Jostling on the other are BEA, IBM, Oracle and Sun, with different versions of technology based on the Java programming language.