Mac to the future

The outer design of future Macs may be known only to Apple's inner sanctum, but we reveal the technologies that will power the next generation of Macs.

WORDS Tony Smith
IMAGE Nils Davey [nils_monkey@hotmail.com]

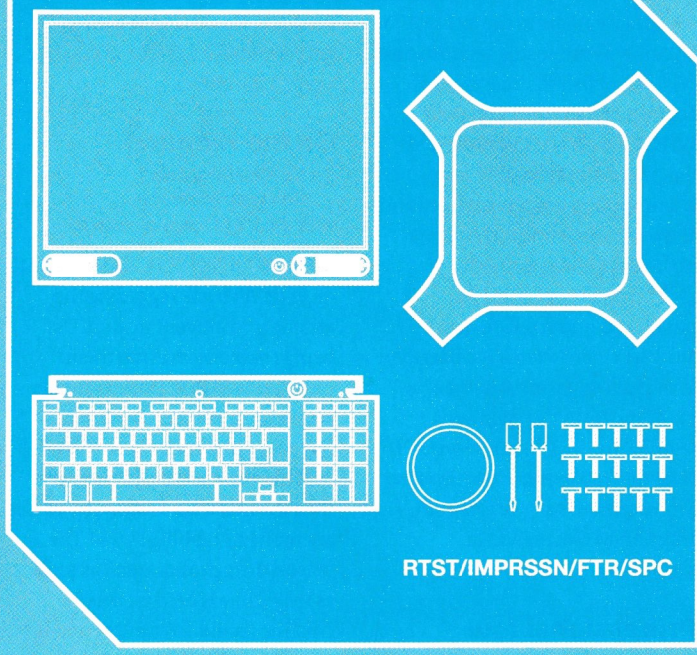

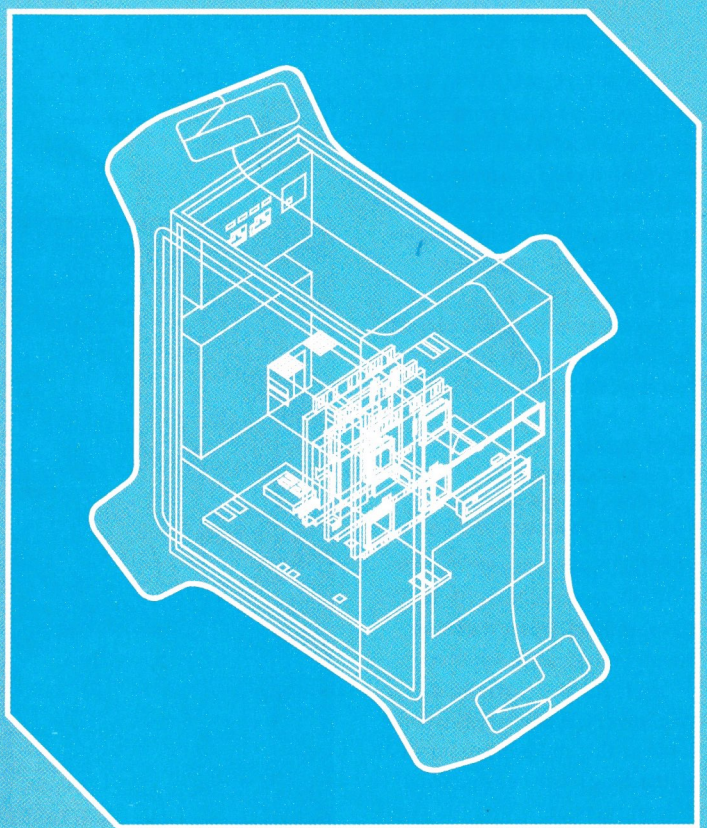

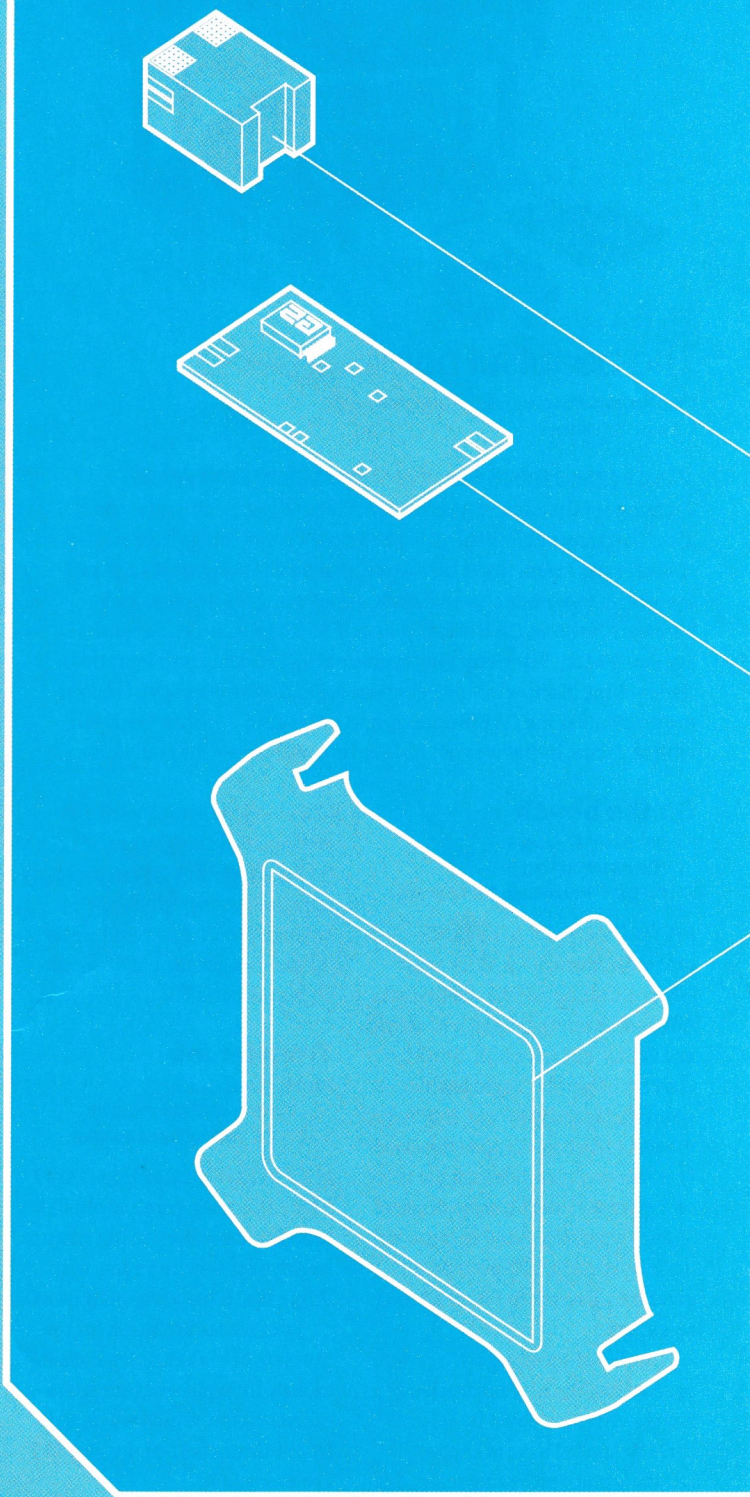

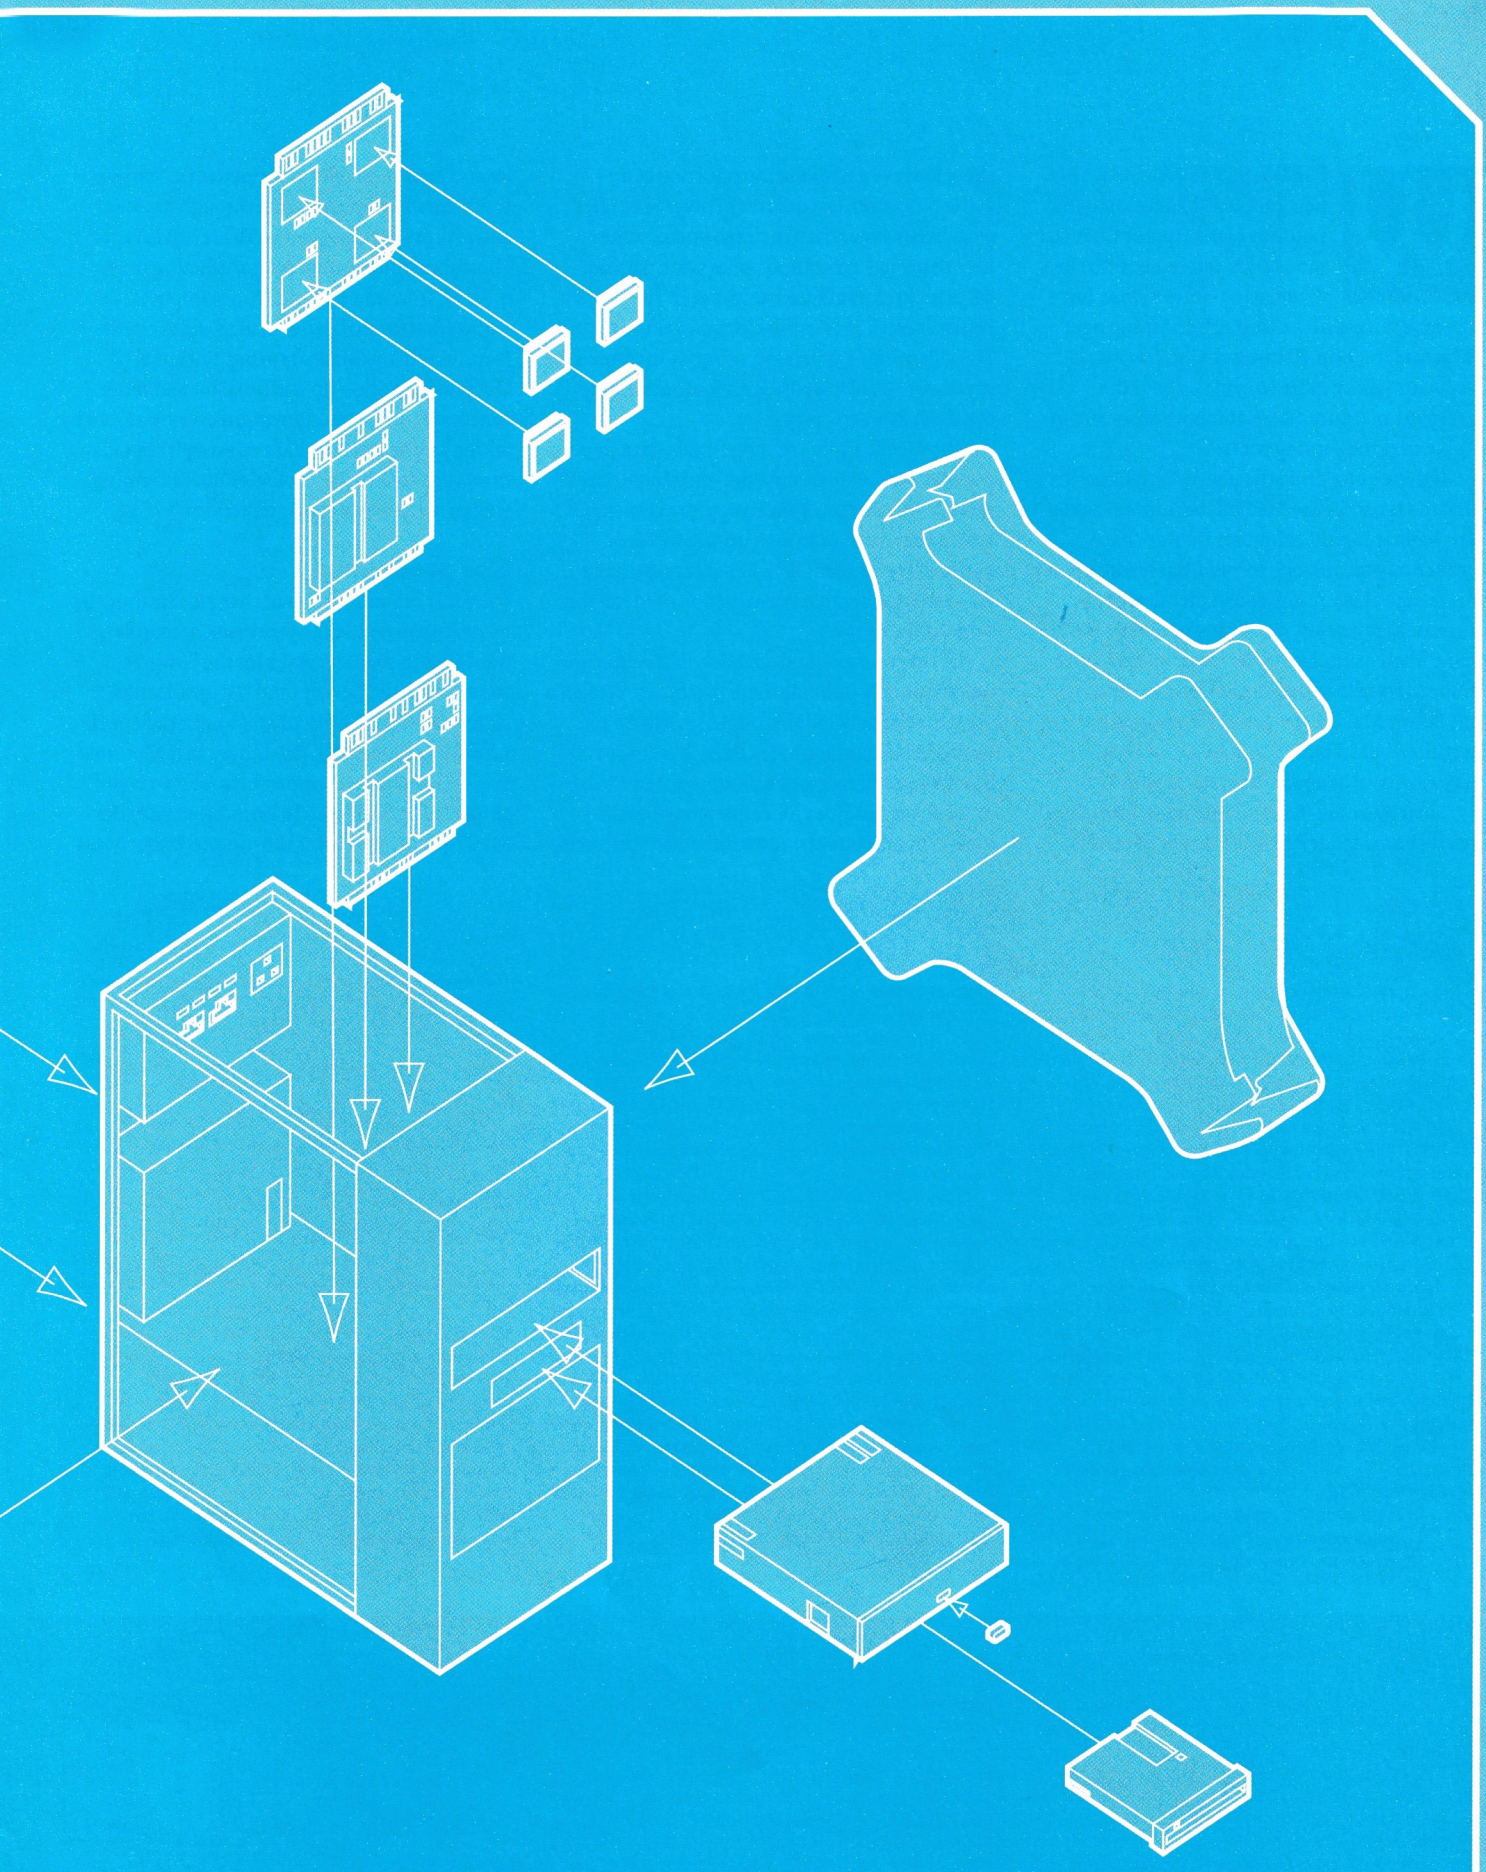

'WE HAVE SOME REMARKABLE new products in development,' said Steve Jobs, Apple's CEO, last month. It's a statement that he's made many times before, but after the release of the iPod, the stunning new iMac and the 17in all-in-one eMac, few Apple-watchers will be quite so cynical about the company in future. So what can we expect from Apple over the coming years?

If there's one thing the new iMac tells us, it's that the future of the Mac will be unpredictable. The hemispherical computer may not have broken any technology records, but it sure tore up the design rules. As the oldest member of the current Mac line-up, the iMac was always the most likely candidate for a radical makeover. But some of its fellows are getting long in the tooth: the Power Mac G4 was given a fresh look last summer, but its basic 'handles' design has remained unchanged since January 1999.

Then there are the relatively poor sales of the professional models. Apple's financial results show that even with the introduction of a 1GHz processor, last quarter's Power Mac sales were no stronger than in the previous three-month period. PowerBook sales were flat, too, despite small upgrades late last year. The PowerBook 600MHz still looks outgunned by a 2GHz Pentium 4 notebook, so Apple increased the processor speed to 800MHz last month using the latest PowerPC processor.

The iBook, Apple's other portable, remains popular. However, despite the introduction of a model with a larger, 14in screen, the consumer machine continues to look underpowered, particularly for Mac OS X. The iMac now has a G4 processor, and buyers may soon demand the same of the iBook.

Of course, there's no question that Apple will eventually upgrade these machines. We can't say what the new models might look like – the unpredictable iMac proved that – but inside the box it's easier. Processor technology is evolving, as are the techniques for connecting it to the system's components, and to other devices, some local, others available on the network. The goal is to accelerate the flow of data through a system and process it more quickly. The result is faster Photoshop filters, more photorealistic games, quicker 3D rendering and iTunes imports.

So what level of performance and range of features will the Macs of spring 2003 offer? Obviously, the line-up will be faster and more functional than now. Apple won't say more, but by taking a close look at what's missing from today's machines and the technologies now emerging, we can make some good predictions.

The processor

The most important area of change, certainly the most anticipated, is the processor. Motorola is tight-lipped about its plans for the PowerPC family, both the G4 class and the upcoming G5 series, reiterating only what it has already made public in its roadmap. This shows that Motorola has delivered on its key promise for G4-class chips: 1GHz clock speeds, introduced with the MPC7455, which drives the top two Power Macs

FASTER CHIPS

Beating the clock

The speed of a Mac's PowerPC processor is set by its clock, a tiny electronic metronome by whose ticks the chip times every operation it performs. The obvious way to increase the chip's speed is to increase the number of ticks the clock makes every second – its clock frequency or clock speed, measured in megahertz and gigahertz (a thousand megahertz).

Alas, it's not that easy. Raising the speed in this way has an unwanted side effect: the chip gets hotter. If it gets too hot, it malfunctions or, in extreme cases, burns up. As clock speeds have risen, Apple's designers have had to devise larger, more complicated cooling systems to keep their processors operational.

Another solution is to shrink the chip. This reduces the amount of heat produced by the chip, because the circuits are smaller, and it gives designers room to incorporate extra circuitry. This is used for features such as the PowerPC G4's Velocity Engine and more on-chip cache memory to hold frequently needed program instructions and data. These facilities can make the chip perform better at a given clock frequency by allowing the chip to do more work for each tick of the clock.

More circuitry, again, means more heat, so designers have to finely balance the chip's size, the complexity of its circuits, the clock frequency they wish to achieve and how much heat it produces.

As chips get smaller, other problems emerge. The chip's circuits are closer together so they begin to interfere with each other's electronic signals, degrading performance. It's harder to drive electrons through smaller circuits, making them less efficient. There are also fewer electrons in the circuit to support the current.

These problems have led to the invention of new materials and techniques. Older chips contain aluminium to connect their circuits; modern ones use copper, which is a better conductor of electricity. The latest generation of PowerPC adds silicon-on-insulator (SOI) technology, which wraps each circuit with better insulation to muffle out interference. Incorporating SOI technology enabled Motorola to boost the PowerPC G4's clock frequency from 667MHz to 1GHz, with just a small rise in the operating temperature of the chip.

Motorola's next trick will be to reduce the size of the chip circuitry, with each transistor it contains reduced from 0.18 microns to 0.13 microns (about one eight-thousandth of a millimetre). This will enable it to boost clock frequencies and/or add more circuitry, such as extra computational units, to allow the chip to perform more tasks per tick.

	Size (mm²)	Transistors (millions)	Process (micron)	Maximum clock speed (MHz)
G2 chip	148	5.1	0.35	233
G3 chip	67	6.35	0.29	400
G4.1 chip	83	10.5	0.18	500
G4.3 chip	106	33	0.18	1000

CHIP SIZES The PowerPC processor has never stopped evolving. The second-generation 604e was a massive 148mm² beast, 40% larger than today's third-generation of G4, yet contained less than a sixth of the number of transistors crammed into the modern chip. As chip elements have shrunk, from 0.35 microns down to 0.29 and now 0.18 microns, Motorola has had more room to add more of them, to make more feature-rich processors, as you can see by the complexity of each chip.

DESIGN ESSENTIALS

Hey, good looking

There's no doubt that Apple is the world leader in the industrial design of the future. While we've yet to see any of the designs Jonathan Ive came up with when studying in Newcastle, perhaps he'll draw some inspiration from these pieces of work, all down by students from Rochester Institute of Technology.

Marcus Conge, who teaches the 3D design course at Rochester, gets his students to design products specifically for Apple as a project. Utilising tools such as Cinema 4D, the students produced the work shown here. 'The exercise is to make something new, but that looks like an Apple product,' says Conge.

Some of the designs have proven to be convincing enough to inspire postings on Mac rumour Web sites, which claim they are genuine Apple prototypes.

While it's not known what Apple thinks of the designs, the students themselves have proved to be good enough to go on to work for many big companies. including, in one case, Apple itself. And Conge has no doubt that Apple will continue to inspire a new generation of product designers.

'Apple represents the most forward thinking for most designers, it is the leader in the shape of its products and the materials it uses,' he says.

Hopefully, the next generation of Mac designs will be equally radical and inspiring.

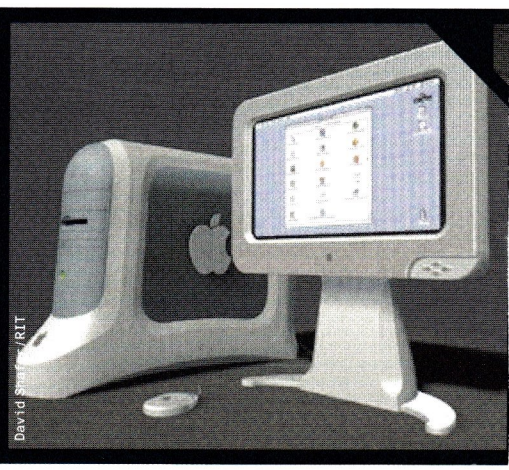

and the 800MHz PowerBook. Last autumn, Motorola subtly updated the roadmapped G4 clock speed from 1GHz to 1GHz+, so it's a good bet that faster successors to the MPC7455 will ship, possibly this summer. In fact, sources at Motorola say the figure of 1.4GHz is most likely.

That's broadly the same timeframe in which Motorola is pledged to offer prerelease 'sample' versions of its first G5-class processor, the MPC8540. Don't expect to see Macs based on this chip – it's aimed at makers of networking products – but it does show the direction in which Motorola will take its desktop processors.

The MPC8540 will be based on Motorola's e500 'core' the bit that does the processing. Alongside it within the chip will be components including Ethernet, Double Data Rate (DDR) memory, PCI-X bus and other I/O controllers, which would ordinarily exist as separate chips. This approach suits the MPC8540's target applications, but not personal computing. The e500 core is founded on a new architecture, and, according to Motorola, is the basis for the G5.

MAC OS ON INTEL

Platform soul

A few years back, it looked as if Motorola would catch up with Intel and offer processors with comparable clock speeds. Intel's drive to 2GHz and beyond put paid to that idea, and while Apple continues to tout the megahertz myth, Intel's growing lead makes the argument less convincing.

So could Apple bite the bullet and adopt Intel or AMD processors? Mac OS X makes it more likely. Its Darwin core already runs on Intel chips, though bringing over other OS X technologies and hardware support software would require further development work.

It is argued that such a move would kill Apple's hardware sales, but an Intel-based Mac wouldn't be the same as an off-the-shelf Windows PC. Apple could easily tailor hardware and software to operate only with each other.

The real obstacle is that software developers would have to convert their applications. Having done this recently for OS X on PowerPC, they may be reluctant to offer and support versions for other platforms. Unless the PowerPC line falls by the wayside – and there's no sign of that – don't expect Apple to switch to Intel soon.

USB 2 AND THE NEXT-GENERATION FIREWIRE

Champing at the bits

Last September, Apple registered the word 'Gigawire' as a trademark. The similarity with FireWire is clear – doubly so when you remember that the successor to the 1394a standard behind FireWire is also known as Gigabit FireWire. Put the words together and you get Gigawire.

It's hard to believe, then, that Apple isn't preparing systems based on the next generation of FireWire, dubbed 1394b, especially given its recent acquisition of FireWire specialist Zayante. This version is not only faster (it extends 1394a's 400Mbits/sec to 800Mbits/sec and provides room to take that quickly to 3.2Gbps) but it allows the use of many more different types of longer cabling.

Although 1394b is in development, it has yet to be fully ratified by the IEEE standards body. However, it is close enough to completion for chip makers, such as Texas Instruments, to develop 1394b interface chips and demonstrate them in public. TI believes we'll see 1394b storage devices ship in the next six months, with 1394b-equipped consumer electronics devices following next year. That gives a broad timetable for Apple to add support for the standard to new Macs.

This is not before time. FireWire's main rival, USB, has already been boosted from 12Mbits/sec to 480Mbits/sec with version 2.0. Market researcher Cahners In-Stat believes USB 2.0 will have completely penetrated the PC market in two to three years time. So while there are few USB 2.0 peripherals now, that's likely to change this year.

Ironically, having driven the mainstream adoption of USB, Apple has fallen behind. The first PC to ship with USB 2.0 support was a Gateway box this year, and Compaq is already building it into some of its notebooks. Apple may well follow suit, possibly when there are sufficient products to justify the move.

USB 2.0 is essentially just a faster version of its predecessor, with which it is compatible. However, the USB On-the-Go project hopes to add support for peer-to-peer connections, allowing devices to exchange data without the need for a host computer. FireWire can do this already. It can also provide devices with power, enabling the iPod to recharge its battery, for instance. That feature can't be adopted by USB without the development of a new connector that might prove incompatible with older devices.

USB 2.0 is likely to limit the need for FireWire devices in the PC arena, but with consumer electronics firmly behind 1394, it will remain a well-supported standard and a key part of Apple's digital hub strategy.

Is the G5 destined for the Mac? Motorola's roadmap shows that the G5 will support a number of technologies, including RapidIO chip-to-chip communications, enhanced multiprocessing capabilities, compact 0.13 micron construction, and clock speeds above 2GHz, that would significantly enhance Mac performance.

Then again, company insiders suggest that the G5 isn't really an evolution of the G4, but a separate product. The latest G4, they say, has also been designed for clock speeds above 2GHz, suggesting Apple can take this processor a lot further than it has already. There's certainly nothing to stop it branding future Macs as G5s, even though they contain G4-class chips.

What we can say is that, G5 or G4, Apple will want to update its top-end Macs regularly, so we can expect faster machines to appear over the next 12 months.

Motorola isn't the only PowerPC producer. IBM is continuing to drive the development of G3-class PowerPCs. The iBook is now the only G3-based Mac, and, as such, Apple probably wants to equip it with a G4. At the moment, it isn't practical and cost-effective to do so, given the G4's higher power consumption and running temperature. However, IBM has started to offer the PowerPC 750FX, a fast G3 that should reach 1GHz by the end of the year. This processor is the logical choice to accelerate the iBook.

Memory matters

Macs may contain the latest PowerPC chip designs, but in many of its components are based on yesterday's technology. Take memory, today's G4s contain memory chips operating at 133MHz. This is the same speed as the system bus, the data channel through which the processor talks to the rest of the computer.

High-end PCs, on the other hand, use significantly more advanced memory. Some manufacturers favour Rambus' RDRAM, others equip systems with DDR SDRAM. Of the two, DDR is expected to become the de facto standard this year, partly because it's a vendor-independent technology, but mostly because it's cheaper.

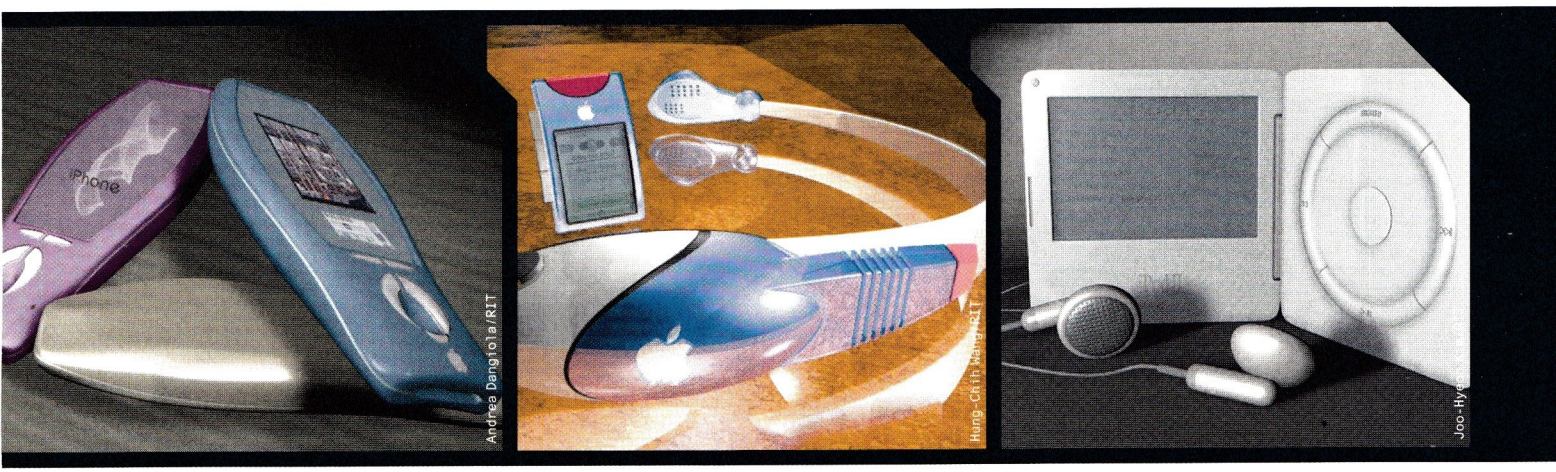

DDR memory is inherently the same as the Single Data Rate (SDR) SDRAM used in today's Macs. Both operate at 133MHz, but clever electronics allow two bits of information to be retrieved from a DDR chip for every single bit read from SDR memory. That makes it appear as if the DDR memory is operating at 266MHz, which is twice as fast as a G4's memory bank. It can keep the processor fed with data far more efficiently than SDR chips can, ensuring less time is wasted while the processor stands idle. Eliminating this 'wait state' becomes more important the higher the processor's clock speed becomes. That's why Intel's Pentium 4 has a 400MHz bus.

Power Macs and PowerBooks already use DDR memory, which is built into their graphics cards and the top-end models' Level 3 cache, because it's more efficient than SDR SDRAM. A DDR main memory bank would significantly improve performance. It would also require a machine's main circuit board to be redesigned: Apple would need a new 'North Bridge' chip, the component that manages communication between the processor, the memory and, in AGP systems, the graphics card. The North Bridge also talks to another chip, the 'South Bridge', which brings together the remaining system components, including the hard disk, AirPort card, Ethernet, USB, FireWire and add-in card controllers. Accelerating the flow of data between these chips and the components they connect is vital to maximising system performance.

HyperTransport

As a member of the HyperTransport Consortium, Apple may turn to this would-be standard to speed up the exchange of information inside the Mac. HyperTransport allows the processor and Bridge chips to talk to each other directly (or 'point to point') and swap information in 'packets', bundles of data not unlike those used to transmit bits and bytes on the Internet.

Nvidia's nForce graphics chipset for PCs is already based on HyperTransport so it is a strong candidate for next-generation Macs. However, there are alternatives to HyperTransport. One is RapidIO, backed by Motorola, which will be built into some, if not all, of the company's PowerPC chips. RapidIO works like HyperTransport but is optimised for networking systems. The bodies behind these two specifications discussed integrating the technologies a while ago, but too much work had been done on each technology.

HyperTransport's development was originally led by chip maker and Intel's arch-rival AMD. Intel promptly began work on a next-generation system bus of its own, known as third-generation input/output (3GIO), a broader technology that not only covers chip-to-chip communication, but also seeks to define a successor to today's AGP and PCI slots. For that reason, 3GIO is backed by the controllers of the PCI standard and will eventually be called PCI Express.

Intel is expected to complete the 3GIO specification by the end of the year. However, it's unlikely to be seen in Macs any time soon, as products based on it won't appear until the second half of 2003. In the meantime, the developers of 3GIO and RapidIO have begun to find ways for their standards to interoperate. AMD's membership of the PCI standard setting body may see it influence 3GIO in HyperTransport's favour. While these moves may not yield a single, unified standard, they should ensure HyperTransport, RapidIO and 3GIO work together, allowing Apple to mix-and-match elements for each or opt for a single specification.

Intel is also working on AGP 3.0, which allows an AGP 8x mode to double the performance of the current standard, AGP 4x, used to drive Power Mac and PowerBook graphics sub-systems. Intel has released initial drafts of the AGP 3.0 specification, but won't say when it will be completed. If the development of 3GIO goes well, there's a chance AGP 8x will never see light of day.

Whatever add-in card standard Apple chooses, it will be able to offer the best graphics technology on the market. Apple courted industry leader Nvidia a few years back, with the result that some of the fastest desktop graphics chips have made it to the Mac ahead of PC versions. Apple has also maintained its relationship with ATI, with both the PowerBook and iBook using ATI chips.

BLUETOOTH

Got no strings...

Apple's release of Bluetooth software last month may be the making of this wireless technology. Despite the hype, Bluetooth has so far failed to become the standard way to connect mobile phones, computers, printers, keyboards and other peripherals as its primary developer, Ericsson, intended it to be.

Bluetooth does for peripherals what AirPort does for networks: it eliminates the cables. Bluetooth operates in the same 2.4GHz band of the radio spectrum as AirPort, hopping throughout the range thousands of times a second to minimise interference and deliver data at 1Mbit/sec.

Unfortunately, by using the same frequencies, it may interfere with AirPort, particularly if their transmitters are placed close together. Apple is unlikely to support something that causes AirPort to malfunction, so the interference problem may be less of an issue than some observers have suggested.

Incompatibilities between different Bluetooth devices have been more problematic. This, and the relatively high cost of Bluetooth communications chips, has hindered the technology's acceptance.

Apple's backing is certainly a boost for Bluetooth, much as it was for USB, FireWire and 802.11b. More important, though, is Microsoft's support. The Windows manufacturer has pledged to support the technology at operating system level in Windows XP this summer. This is likely to encourage other peripheral makers to sign up. A number of analysts from the likes of Frost & Sullivan and Gartner are claiming that sales of Bluetooth communications chips will be well up this year, leading to the wide availability of Bluetooth-based products in 2003.

As a technology leader, Apple may well want to stay ahead of the curve. Indeed, Apple recently dropped the infrared IrDA port from the PowerBook G4 line, and is promoting Bluetooth for PC-to-phone or PC-to-PDA communications instead. Building the technology into new Macs can't be far away.

The company is already working on the software to integrate Bluetooth into Mac OS X. For example, version OS X 10.2 allows users to dial a number in the Address Book application using a Bluetooth-equipped telephone.

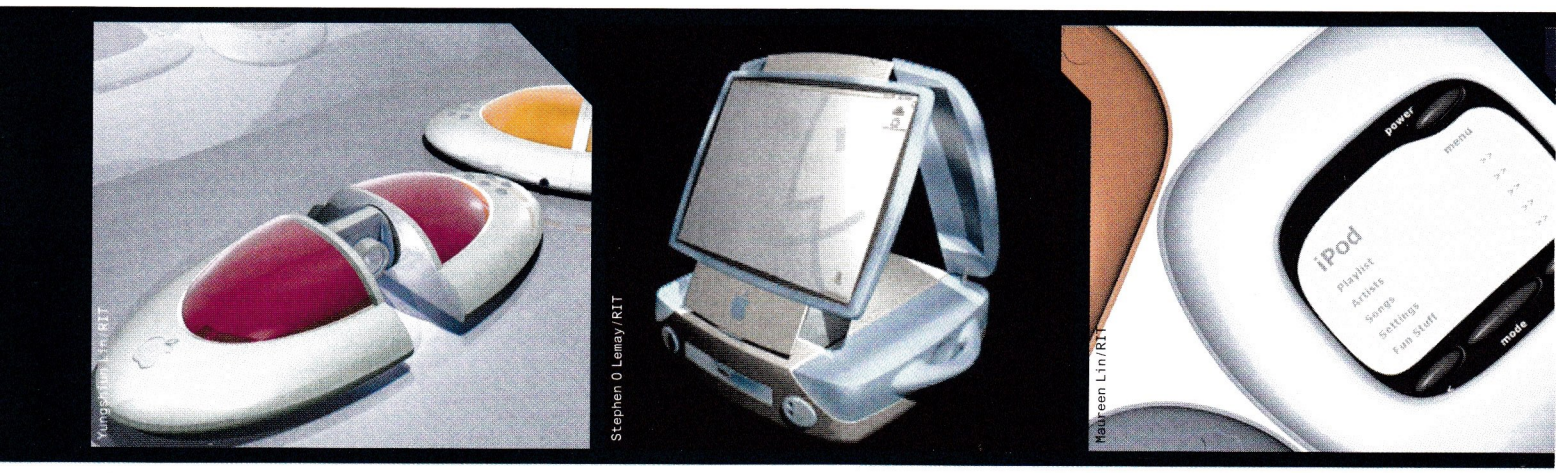

WIRELESS NETWORKING

Easy as a, b, c

AirPort, Apple's wireless networking system launched in 1999, helped kick off a mobile computing revolution. The technology on which it's based, the 802.11b open standard, had been around for some time, but Apple brought it into the mainstream.

Today, there are thousands of 802.11b networks operating in homes and offices in the UK and the US. Every modern Mac supports it, and PC owners can play too, with a variety of 802.11b products from companies such as Proxim, Intel, 3Com and Cisco. Dell, IBM, Compaq and Toshiba are following Apple's lead and offering the technology in their notebooks.

AirPort has its limits, though. It operates in the 2.4GHz band of the radio spectrum, which it shares with noise from microwave ovens and other sources, and it's optimised to operate with a direct line of sight between base station and Mac. Together with complex error-correction technology, these factors lower its data-transfer rate from a potential maximum of 11Mbits/sec to less than half of that. Mac owners used to sending massive Photoshop files quickly over a wired network may find AirPort's performance disappointing.

Fortunately, improvements are being made to two standards, 802.11a and 802.11g. Both raise the performance of 802.11b to 54Mbits/sec and don't demand a good line of sight to operate efficiently. However, while 802.11g operates in the 2.4GHz band for compatibility with 802.11b, it's susceptible to the same noise. 802.11a operates in the noise-free 5GHz band, but that renders it incompatible with 802.11b. This might make 802.11g the best choice, but 802.11a is here now and 802.11g has yet to be ratified as a standard.

That, says Anthony Fulgoni, Proxim UK's chief, won't happen before the end of the year, by which time 802.11a, already available in the US, will be released here. Authorisation to use the technology anywhere in Europe is expected to be granted by the EU this summer.

Fulgoni admits that 802.11a equipment is more expensive than 802.11b, but the difference is small, and he expects businesses to begin installing the faster product by the autumn. This will be alongside existing 802.11b networks to allow the older technology to remain operational. He believes the rapid and widespread take-up of 802.11a (and the fact it's only slightly more expensive than 802.11b) will end 802.11g before it's even available.

With 802.11a approved in Europe and the US, it's a good bet that Apple will update AirPort some time in the next 12 months.

Graphics chip development, like that of computer processors, is driven by the need to increase clock speeds and add more complex circuitry. Both of these aspects are enabled by the production of more compact chips. Graphics processors are now made using the same processes as CPUs, allowing Nvidia and ATI to build in more sophisticated rendering technologies, such as the programmable pixel shaders, seen in the latest Radeon 8500 and GeForce4 chips.

That programmability will be enhanced over coming generations of these chips. The arrival of 64-bit colour will also enable some staggering lighting and texture effects. Today, a pixel's colour is defined by 32 bits. This is equal to eight bits each for the three primary monitor colours (red, green and blue) plus eight to determine the pixel's transparency. That means a pixel's redness, for instance, can be defined only on a scale of 0 to 255. That's enough to show a photorealistic image, but start mixing in the effect of multiple light sources, and it soon proves insufficient.

With such small numbers, you get rounding errors, which tend to wash out the colour. By defining each red, green and blue colour component with 16-bits each, blending colours becomes far more subtle and more accurate. Our eyes won't see any more colours, but the image will be richer and more realistic. Of course, all these extra bits will require more memory to store them, so graphics cards of the future will have more RAM than most Macs currently contain.

As will the Macs themselves. More memory, bigger hard drives, faster communications, fewer cables, a greater reliance on wireless technology, multiple multigigahertz processors: all these features and more will certainly shape this year's Macs and those coming further down the line, as they always have. But it's Apple's capacity for innovation in other areas that will define the Mac platform and keep it ahead of the pack.

There's one thing we can be sure of, though: future Macs will continue to be, in the words of Steve Jobs, 'remarkable'.

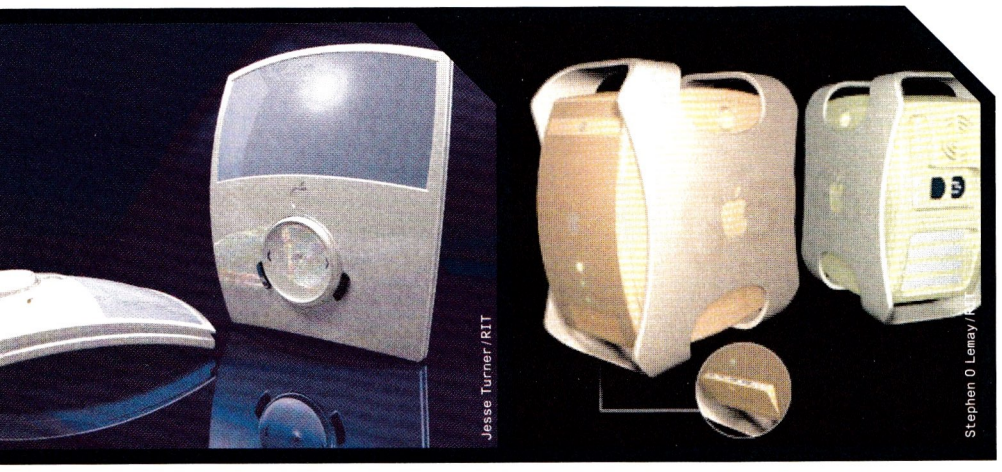

JARGON

Talk the talk

■ AGP
Accelerated Graphics Port. A technique for connecting a graphics card to the processor, bypassing the busier (in other words, slower) system bus.

■ Bluetooth
A wireless technology for connecting peripherals to each other. It currently supports a 1Mbit/sec data transfer rate, so it's not for high-bandwidth applications but for connecting devices such as keyboards, mice, graphics tablets and synching PDAs.

■ Bus
A channel through which data is transmitted. A Mac contains a number of buses, such as the system bus, which links the processor to memory and the input/output system, the FireWire bus, and the ATA bus, which connects the hard drive to the system.

■ Cache
A small bank of local memory that records frequently accessed information. The cache is kept close to where the data is needed, usually the processor, to reduce access times.

■ Chipset
The collection of chips that links the processor to other system components, such as memory and add-in cards, and manages the flow of data within the computer.

■ Clock speed
Also known as clock frequency, this is a measure of how many times a second processor's 'heart' beats. It is measured in megahertz (millions of beats a second) or gigahertz (billions of beats a second).

■ IEEE 1394
The official name of the communications standard better known as Apple's FireWire or Sony's iLink. 1394a permits data to be transferred at 400Mbits/sec over 4.5m shielded cables. 1394b extends those figures to 3.2Gbits/sec over 100m cables of various types.

■ IEEE 802.11b
The official name of the wireless networking standard better known to Mac users as AirPort. It allows data to be transmitted at 11Mbits/sec through 11 channels in the 2.4GHz band of the radio spectrum.

■ Instruction Set Architecture (ISA)
Different makes of processor have different instruction sets, which is why they're usually incompatible. PowerPC instructions won't make any sense to an Intel chip, but since Intel and AMD chips both use the x86 ISA, they are compatible.

■ Micron
The measurement of the size of a key circuit element – the 'gate' – within a processor and the process by which the chip is made. Today's PowerPC has a 0.18-micron chip: the next generation will be 0.13-micron parts.

■ SDRAM
Synchronous Dynamic Random Access Memory. The favoured design of computer memory, capable of being synchronised with the Mac's processor to improve the flow of data between them.

■ USB
Universal Serial Bus. The commonly supported version 1.1 of this peripherals interface standard supports speeds of up to 12Mbits/sec. The more advanced but still uncommon version 2.0 raises this to 480Mbits/sec.