The Economist

Touching the future

Computing: Touch screens are becoming an increasingly popular way to control mobile phones and other devices. How does the technology work, and where is it heading?

THE proliferation of touch screens in electronic devices over the past two or three years has been so rapid that you may have found yourself trying to press an on-screen button or icon when sitting at your computer only to realise, much to your frustration, that it is not a touch screen. Many mobile phones, most famously Apple's iPhone, now have touch-screen interfaces, as do satellite-navigation systems and portable games consoles. Confusingly, however, most computers do not—so far.

But that may be about to change. Microsoft has already demonstrated a prototype of Windows 7, the next version of its flagship operating system, based around "multi-touch" capabilities, which allow a touch screen to sense more than one finger at once. As well as being able to press buttons, tap icons, call up menus and scroll windows, users will be able to rotate and stretch on-screen objects using two fingers at a time, as they already can on the iPhone. For its part, Apple is rumoured to be working on new versions of its desktop and laptop computers with touch screens. It has already taken a half-step in this direction by putting multi-touch trackpads into its laptop computers.

So the touch screen could be on the verge of becoming a standard part of computer interfaces, just as the mouse did in the 1980s. Many people thought that would never happen: surely switching between keyboard and mouse would slow people down and make them less productive? In fact, mouse-driven interfaces can be far more efficient, at least for some tasks. The same seems likely to be true of touch-screen interfaces. The touch screen will probably not replace the mouse and keyboard, but will end up being used for some tasks.

How touching

Like the mouse, the touch screen has been around for quite a while. Indeed, one of the ironies of Microsoft's demonstration of Windows 7 was that it was carried out using a screen made by Tyco Electronics Elo TouchSystems, which pioneered the commercialisation of touch screens in the 1970s. Today countless supermarket checkouts, restaurant tills, automated-teller machines, airport check-in kiosks, museum information-booths and voting kiosks use touch screens. Intuitive and user-friendly they may be, but new they are not.

Elo TouchSystems' first touch screen grew out of the work of Sam Hurst, a physicist at Oak Ridge National Laboratory, who was on teaching leave at Kentucky University. His original touch sensor had nothing to do with screens. It was designed to save time when manually inputting a huge stack of spectrometry readings—in essence, a series of jagged lines on graph paper—into a computer. Rather than waiting two months while graduate students did the job by hand, Dr Hurst decided that there had to be an easier way.

What he came up with, in 1971, was a device consisting of two pieces of electrically conductive paper, separated by an insulating layer. A needle poked into the paper acted as a bridge between the two layers. Dr Hurst could determine the needle's horizontal and vertical position by applying a voltage across the sheets and using voltmeters to measure the change in voltage when he inserted the needle. He could then place a spectrometry plot on the device, trace over its jagged line using the needle, and read off the co-ordinates quickly from the voltmeters as he went along.

Touch screens may be intuitive and user-friendly, but they are not new.

Dr Hurst realised he had hit on something, but it took him a while to work out precisely what. He founded Elographics (now called Elo TouchSystems) with nine colleagues and tried selling his device. "It turned out it wasn't a good business plan, because we were aiming it at a scientific market," says Dr Hurst, who is now retired. It was not until 1973 that he thought about adapting his invention so that it could be laid over screens.

Elographics then did away with the needles and developed a new device based on two transparent, conductive surfaces separated by a small air gap maintained by a series of spacers. When a finger pressed the two surfaces together, they made contact, and the finger's position could be determined using the technique that identified the position of the needle in Dr Hurst's original device. The result was the first "resistive" touch screen. It is used in many devices and is still the commonest type of touch screen today, accounting for 91% of unit sales in 2007. (This type of screen is used in the Nintendo DS games console, for example.)

But this is not the only way to build a touch screen, and nor was it the first. Bill Buxton, an adjunct professor of computer design and interaction at the University of Toronto and a principal researcher at Microsoft Research, says IBM was working on infra-red touch screens as early as 1965. These are based on a grid of invisible infra-red beams that criss-cross the surface of a screen. When a finger touches the screen, it breaks a horizontal and a vertical beam, which reveals its position.

This kind of touch screen was used in a computer-assisted instruction system, called PLATO IV, that was developed in the late 1960s at the University of Illinois at Urbana-Champaign, and released in 1972. It allowed students to answer questions by touching the screen, and was a genuine "touch" screen, because there was no need to apply pressure to the screen. "Surface acoustic-wave" touch screens work in a similar way, with beams of ultrasound rather than infra-red light.

The iPhone uses yet another technology, called a capacitive touch screen, which has even more distant and unusual origins. As far back as 1953 Hugh Le Caine, a Canadian pioneer of electronic music, was developing capacitive sensors for an early synthesiser, called the Sackbut, in which the position of his fingers controlled the timbre of the sound.

A capacitive touch screen is covered with a transparent conductive coating, to which a voltage is applied, establishing an electric field across the screen. When a finger touches the screen it disrupts this field, as the human body's natural capacitance causes a local build-up of electric charge. Electrodes around the edges of the screen sense the change in charge distribution, allowing the finger's position to be determined (see diagram).

In fact, the iPhone uses a more elaborate form of capacitive touch screen. A grid of electrodes divides its display into an array of hundreds of small, separate touch elements. Each of these can be triggered independently, enabling the iPhone to sense more than one finger at once, and making possible "multi-touch" features, such as pinching an image to shrink it. Multi-touch is also an old idea: the first prototype was cooked up in 1984 by Dr Buxton, but it was a touch pad, since it was not built into a screen. "The reason we didn't make them display devices is not that we didn't want to—it was because we couldn't figure out how to make them cheap enough to put on a screen," he says. For that he credits Bell Labs, which came up with the materials necessary to make a capacitive multi-touch device that could be combined with a computer display.

You can touch this

If touch screens have been around for so long, why did they not take off sooner? The answer is that they did take off, at least in some markets, such as point-of-sale equipment, public kiosks, and so on. In these situations, touch screens have many advantages over other input methods. That they do not allow rapid typing does not matter; it is more important that they are hard-wearing, weatherproof and simple to use. "In a museum, if you can't figure out how to use an input device, you aren't going to use it," says Joel Kent, a senior scientist at Elo TouchSystems.

But breaking into the consumer market was a different matter entirely. Some personal digital assistants, or PDAs, such as the Palm Pilot, had touch screens. But they had little appeal beyond a dedicated band of early adopters, and the PDA market has since been overshadowed by the rise of advanced mobile phones that offer similar functions, combined with communications. Furthermore, early PDAs did not make elegant use of the touch-screen interface, says Dr Buxton. "When there was a touch interaction, it wasn't beautiful," he says.

That is why the iPhone matters: its use of the touch screen is seamless, intuitive and visually appealing. When scrolling quickly through lists, for example, the lists keep moving, apparently under their own momentum. On-screen objects behave in physically realistic ways. "Apple did an outstanding job with the iPhone in every aspect," says Dr Buxton. "People are starting to demand the same quality of design: they don't just want function."

Until recently, the computing power and graphics capabilities of desktop computers, let alone hand-held devices, were not good enough for elegant touch-screen interfaces to work, says Jeff Han, founder of Perceptive Pixel, a touch-screen spin-out from New York University's Courant Institute of Mathematical Sciences. And even if they had been sufficient, the public might not have been ready for such interfaces. "In the 1990s people were still getting used to Windows 95," he says. "But right now the public is ready—even the most lay person can use a mouse."

Fingers crossed

How touch screens work

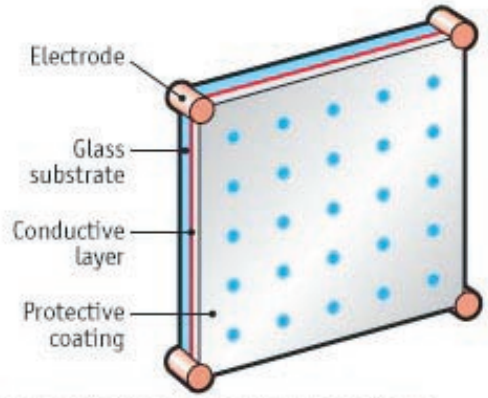

A capacitive touch screen is covered with a transparent conductive coating, to which a voltage is applied, establishing an electric field across the screen.

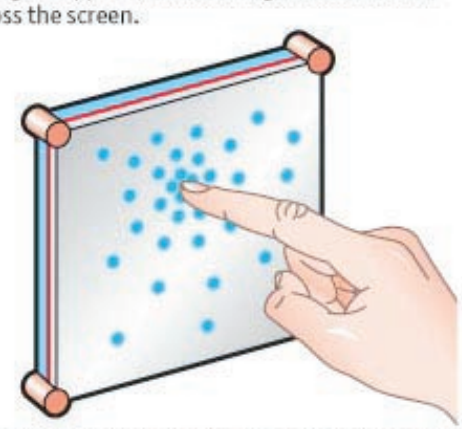

This field is disrupted when a finger is placed near the screen, as the human body's natural capacitance (ability to store electric charge) causes a local build-up of electric charge. Electrodes around the edges of the screen sense the change in charge distribution, allowing the finger's position to be determined.

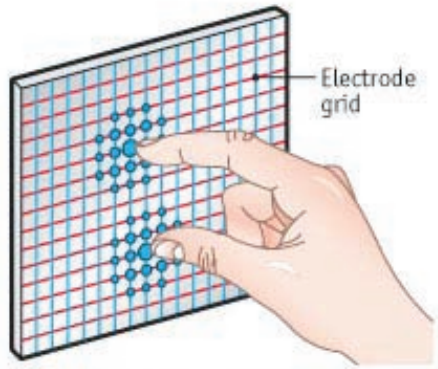

Apple's iPhone uses a more elaborate form of capacitive touch screen, with a grid of electrodes that create hundreds of small, separate touch elements. Each of these can be triggered independently, allowing the iPhone to sense more than one finger at once, and making possible "multi-touch" features such as pinching an image to shrink it.

Source: *The Economist*

Microsoft and Apple

Another factor that has held back touch screens is a lack of support for the technology in operating systems. This is a particular problem for multi-touch interfaces. Modern operating systems, driven by keyboards and mice, are unable to cope with a system that is, in effect, like connecting several mice at once. Instead, they are based on the idea of a single cursor that glides from one place to another. In developing his touch screens, Dr Han says he was forced to create his own operating system. "We didn't want to," he says. "But Windows and Linux really don't understand more than a single point."

Apple and Microsoft have now cottoned on, however, and have added multi-touch support to OS X (the operating system that powers Macs and iPhones) and Windows. It is still not clear what features Windows 7 will include, but some clues can be gleaned from Microsoft's Surface, a computer built into a tabletop with multi-touch input. It is an expensive device that is initially aimed at the business market: it is being deployed in hotels, restaurants and casinos.

Surface uses under-the-table projection to produce the tabletop image. The screen is also illuminated from below with infra-red light, which reflects off objects in contact with the screen. These reflections are detected by cameras, also mounted beneath the table. "The very first one had a single camera lying on the floor, looking up at a savaged IKEA table," says Andy Wilson of Microsoft Research, who started developing Surface with his colleague Steve Bathiche in 2001.

The beauty of using cameras in this way is that as well as detecting multiple touch inputs, and supporting several users at once, they can also recognise objects placed on the table's surface. For example, users of the system in a restaurant could place empty glasses on an on-screen icon to indicate that they would like a refill.

Perceptive Pixel's huge touch screens also use cameras to detect touch inputs, but in a different way. The image on the screen is back-projected, but the touch-detection relies on a technique called "frustrated total internal reflection". A transparent acrylic sheet is illuminated by a number of infra-red light sources placed around its edges. The infra-red light is confined within the acrylic sheet, because beams of light are internally reflected off its front and back faces. When someone presses the sheet with a finger or stylus, however, this "total internal reflection" is disrupted, or frustrated, and some of the light leaks out of the sheet at the point of contact. This light is picked up as a bright spot by an infra-red camera behind the screen. Touching the screen in many places at once results in several spots.

Inappropriate gestures

Dr Han is a virtuoso when it comes to demonstrating the potential of touch screens. He calls up images and charts and navigates through three-dimensional maps, panning and tilting them using special multiple-finger gestures. He has also developed special gestures to call up menus and perform other commands. When you have a very large screen, covering an entire wall, you do not want to have to walk to one corner to activate the Start Menu, he says. So drawing a loop with one finger might call up the options menu, for example. Designing such gestures can be tricky. "You want something that's hard to do accidentally, but not hard to learn," he says.

Microsoft is also developing gestures, and Apple has already introduced several of its own on its multi-touch enabled laptops, such as two-fingered dragging to scroll, and three-fingered flicking to go forward or back a page in a web browser. The danger is that a plethora of different standards will emerge, and that particular gestures will mean different things to different devices. Ultimately, however, some common rules will probably emerge, as happened with mouse-based interfaces. "Double-clicking didn't used to be universal," notes Dr Buxton, but now it is accepted as the standard way to open a program or document on most computers.

The double click does not translate terribly well to touch screens, however. This has led some researchers to look for alternatives. In developing his multi-touch screen, Dr Han found that there can be more to touch input than simply detecting contact. He has found a way to determine how much pressure is being applied. Adding a thin polymer layer, scored with microscopic ridges, to his touch screens causes the bright spot created by a finger touching the screen to vary in size and brightness depending on the pressure. This makes it possible, for example, to drag an item on the screen and then, by pushing harder, to slide it under another item.

Quite what will be done with multi-touch and pressure-sensitive screens is still unclear, says Lorna Wood, marketing manager of Elo TouchSystems. "A lot of the applications have yet to be developed that really take advantage of this technology," she says. But touch screens seem likely to become more widespread in desktop PCs, laptops and mobile phones. According to iSuppli, a market-research firm, sales of touch screens will increase from 341m in 2008 to 833m in 2013.

Despite the iPhone's success, it may prove to be PCs, rather than hand-helds, that benefit the most from touch-screen technology. That is because touch screens, like mice, are best suited to manipulating information, rather than inputting it in the first place—an area in which keyboards remain unchallenged. PCs with keyboards and touch screens (not to mention mice or trackpads too) could offer the most flexibility, letting users choose the appropriate input method for each task. As Dr Buxton puts it: "Everything is best for something, and worst for something else."

elo TOUCHSYSTEMS

Tyco Electronics
Our commitment. Your advantage.

Elo TouchSystems ı 301 Constitution Dr. ı Menlo Park, CA 94025
Phone: 1-800-557-1458 ı customerservice@elotouch.com ı www.elotouch.com