System Security for Cyborgs

Ross Anderson
Cambridge University Computer Laboratory
Ross.Anderson@cl.cam .ac.uk

Within our lifetimes, we will carry dozens of smart devices about with us. Many will be worn, but some may be embedded. Many will just be tools, like today's phones and iPods, but there will be medical monitoring devices too, and some people will have things like pacemakers on which they seriously depend. We may come to depend on the tools too. Human beings may be the ultimate 'embedded systems'.

Devices will interact with each other. Communications between consumer devices such as phones, cameras, GPS receivers and MP3 players are already standard. Researchers are working on more subtle ideas, such as using a heart-rate monitor to stop your mobile phone ringing when you're cycling. There will be many more.

Over time, there will be safety and security scares. Science fiction writers have come up with many lurid ideas, from assassination by malware to the Nam-shub of Enki. The intrusion of software industry practices such as platform monopoly, product tying and region-coding into such an intensely personal domain as the human body may also lead to interesting tussles.

In this talk, I'm going to sketch on a narrower canvas – the security of devices that are worn, or implanted, for medical or health reasons. By security I mean that systems remain dependable in the face of malice, error or mischance. Dependability in the medical context will mostly mean safety, but from time to time also confidentiality.

System safety problems with medical monitoring systems are not new. During the anxious build-up to the year 2000, government consulted about the reactivation – in the event of disaster – of the 'telephone preference scheme', a Cold War civil-defence system whereby only important people such as policemen and GPs would be able to initiate phone calls. The rest of us would have to wait to be called. The response from general medical practice was that thousands of people now depend on remote monitoring for their care, and if they could not be put on the preference list then they should be hospitalised for the duration of the anticipated emergency.

Tensions between safety and security are also becoming familiar. When a medical device is based on a Windows platform, and Microsoft issues a critical security patch, do you apply it or not? If you don't, your medical device might be hacked and turned into a spambot. If you do, then your machine is no longer in a configuration approved by the regulator, and your malpractice cover could be at risk [1].

As the health and well-being of more and more people come to depend on worn or implanted devices, and on the back-end systems with which they communicate, the dependability of these devices and of their supporting infrastructure will become more critical. How can we ensure that communications between a patient's medical devices are not jammed (say by a mobile phone) or confused (for example, if I sit down on a bus next to another patient with the same type of device, and his sensor signals are wrongly acquired by my equipment)? Where can we look for guidance about what's likely to go wrong and what we might do about it?

The experience of the motor vehicle industry might be worth a look.

A late-model luxury car may have 50-100 communicating computers on board, running all sorts of services from the obviously critical, such as brakes and traction control, through air-conditioning and seat positioning, to services such as entertainment and navigation. It is about as close as civilian equipment gets to the 'pervasive computing' dream of computers embedded invisibly everywhere, unobtrusively providing assistance in many small ways, with serious thought having been given to integration and usability.

Just as government had not realised in 2000 that the availability of the telephone system had become medically critical, few people had relaised that a car's window motors, door locks and air conditioning had become safety-critical. Then, in May 2003, the Thai finance minister, while riding in a BMW through Bangkok traffic to a conference, experienced a software failure that stalled the engine, locked the doors, jammed the windows shut and turned off the airconditioning. Had a nearby hotel doorman not brought a hammer and broken a window, the minister and his chauffeur would have suffocated [2].

Until very recently, information security mechanisms were added to vehicles piecemeal, along with each new feature such as remote locking, road toll tags, and stolen vehicle tracking. But now manufacturers are starting to find systems interacting with each other. As customers demand more and more features, the problem is getting worse, and as the subsystems come from a large number of different suppliers, system integration is becoming one of the biggest quality headaches for the car industry. It particularly affects expensive models, as these adopt the latest gadgets early.

Now, IBM has launched 'Embedded Systems Lifecycle Management', a new business venture aiming to integrate and manage automotive software and electronics. They claim that 32 percent of warranty costs are due to dealer service visits at which no problem is found – often relating to transient faults that arise when one 'smart' device interferes with another. Failures can force the driver to stop the car and restart it; they can cause false alarms; and diasgnosing them is so hard that attempts to fix them can involve replacing numerous sensors [3].

Usability is also a growing issue for vehicle electronics, as it is for computer systems generally. Often customers simply cannot figure out how to perform some necessary task because their model of the world differs from that used by the designers of the interfaces, or because there are just so many menus that they give up before finding the right one. It may take several attempts for a car maker to produce controls with which its customers are really comfortable. However, once a platform can be programmed, there is constant pressure to add new features, as the marginal cost of this is low.

Growing evidence suggests that embedded systems tend to acquire features until they reach the point of frustration – when the utility gain from adding one new feature equals the utility loss resulting from confusion and from being unable to use the existing features well [4]. The net effect of usability research is not more usable systems, but equally frustrating systems that frustrate us at a higher level of complexity. Perhaps user tolerance of real or perceived 'bugs' will be lower for embedded devices than for cars or mobile phones; we will have to wait and see.

Usability is related to risk, as operational failures bring accidents in their wake. Just as we currently accept several thousand deaths per annum from road traffic accidents, because of the usefulness of the motor car, so our grandchildren may well accept some deaths from feature interactions between implants and prostheses, and from the effects of faulty monitoring equipment. But the culture clash between computer-industry attitudes to dependability and the more cautious attitudes of physicians may be even sharper, over the next thirty years, than the clash between computer people and telephone-industry people over the last thirty.

Maintainability brings another set of challenges. However expensive it may be for car makers to retrofit vehicles with bugfixes and interface improvements, it will be worse for medical device vendors. Sensor replacement may involve invasive surgical procedures rather than a simple garage visit. So thought must be given to mechanisms for remote software upgrade. This is harder than it looks! Not only must upgrades be properly authenticated and authorised, but it must also be difficult for the upgrade mechanisms to be abused so as to deny service (for example, by uploading a genuine copy of new software but halting the upload halfway through).

The use of commercial-off-the-shelf components can greatly reduce costs, and in some applications one has little choice. For example, the costs of developing new mobile phone platforms are so high that almost all specialist GSM / GPRS devices are adapted from standard mobile phones. But this can lead to platform security issues and to lifecycle issues in embedded systems. For example, some high-end German cars come with an embedded T39 mobile phone for remote software upgrade. This model is already obsolete, yet these cars may be on the road for 20 years or more. In addition, a number of mobile-phone platforms are starting to experience attacks by computer viruses and worms, and we have seen the first reports of cars being affected [5]. And while PC software vendors are used to shipping security patches, mobile phone vendors are not – they hope you will throw your phone away after a year and buy a new one. There are some interesting tensions developing here.

Privacy issues may arise in several ways. First, devices may leak information locally; if body-area networks carry medical data, then the presence of a particular medical device may be detectable nearby. The world's intelligence services have long tried to collect personal health information on heads of government and their likely successors, while journalists show even more interest in other types of celebrity.

Second, information may leak centrally. Over the last decade or so, the NHS has been collecting summary information on all episodes of secondary care (via the Clearing and HES systems) and a similar data-collection exercise is now underway for GP data. Unless this trend is reversed, I expect that worn and implanted medical monitoring devices will generate increasing quantities of data, which will be made available to a growing number of applications in research, clinical audit, administration and marketing. It is a matter of wonder to people aware of this centralisation of personal health information that neither the data protection authorities nor the public have yet become alarmed about it. Past some threshold, or following some scandal, that may change radically [6].

The next bundle of issues concerns market structure and competition policy. Markets in information goods and services tend to monopoly, because of the combination of high fixed and low marginal costs, high switching costs, and the presence of network effects. The value of many IT companies depends on how tightly they can lock in their customers, and profits are often maximised by strategies that involve bundling and tying products. We commonly find information security mechanisms used to support such business models – a topical example being the authentication chips that prevent a computer printer from using ink cartridges supplied by a third party [7]. Cheap printers are subsidised by expensive ink cartridges. Although economically efficient (as home and business users can be fed off the same production line), this is unpopular, and has led to intervention by the European Parliament [8].

As more and more industries incorporate computers and digital communications into their products, so they will come to look more like the software industry. There will be the good (flexibility), the bad (complexity), and the ugly (monopolies). There appears to be nothing about the medical device business to push it in a different direction; indeed, surgical implantation greatly increases the costs of switching from one make of electronic device to another. Healthcare industry regulators may end up having to study the history of the antitrust cases against AT&T, IBM and Microsoft, and become as intimately familiar with the issues surrounding open source software and cryptographic lock-in, as they currently are with pharmaceutical patent licensing.

We can also expect that, as information security mechanisms support goals that benefit different stakeholders – such as safety, privacy, cost control and lock-in – they will end up serving more than one master. The patient will ask: 'is this a safety system which helps me, or a control system which restricts me?' [9] The resulting conflicts of interest may stress the relationships between patients, physicians, vendors, insurers and regulators.

There may be new debates about the limits of government power. It is already policy that all motor vehicles in the UK will carry GPS-based road-tolling equipment by 2010, which will give the police and intelligence services a location history of every car in the country. If a similar location history of individuals becomes technically possible, will the agencies demand it and will Parliament grant it? It is also policy that future vehicle systems will enable police to slow down and stop all vehicles in an area, and to immobilise cars whose tax or insurance is over-due for renewal. It would clearly be convenient for the police to immobilise wanted persons remotely; will they be allowed to? What about the patient who's forgotten to pay the software license renewal on her pacemaker? On the memory prosthesis embedded in her skull? And what if some of the space in the prosthesis is used by an MP3 player – if she cancels her music subscription, will the DRM service delete the associated memories? If so, will the medical ethics people prevail and ensure the amnesia is gentle, or will the marketing people get the upper hand and leave her with an unbearable sense of loss?

Finally, there are sordid implementation details. If the experience of the motor-vehicle electronics industry is repeated by the medical device manufacturers, they will waste a lot of money by ignoring the existing expertise on information security engineering and reinventing everything for themselves. Car makers devised their own encryption algorithms, their own authentication protocols, and their own information flow controls. Often they got it wrong, and are only now starting to get in experts to put it right. Cars started using cryptography in remote key-entry systems in 1992, but the first conference on electronic security in cars did not happen until 2003 [10].

To sum up: if our grandchildren are going to be cyborgs, their design will pose many fascinating challenges for the information security engineer. These will range from the profound matters of public policy to tricky technical details. Among the hardest medium-term issues, though, are likely to be those related to safety. If I were advising a gifted security research student who was looking for a thesis topic in this general area, I might well suggest usability and maintainability as the most important – and the most challenging – of the problems.

Acknowledgement: I am grateful for discussions with members of the security group and with Karen Spärck-Jones. This article is licensed under the GFDL; its residual copyright will subsequently be assigned to the IEE.

REFERENCES

[1] Ellen Messmer, "Fed up hospitals defy patching rules", *Network World* 9 August 2004, at http://www.nwfusion.com/news/2004/080904patchfights.html

[2] Julia Scheeres, "Teched-Out Cars Bug Drivers", *Wired*, 29 June 2004, at http://www.wired.com/news/print/0,1294,63846,00.html

[3] Tim Moran, "What's Bugging the High-Tech Car?". *New York Times*, 6 February 2005, at http://www.nytimes.com/2005/02/06/automobiles/06AUTO.html?oref=login

[4] Andrew Odlyzko, "The Visible Problems of the Invisible Computer: A Skeptical Look at Information Appliances", in *First Monday* v 4 no 9 (1999) at http://www.firstmonday.org/issues/issue4_9/odlyzko/

[5] Dan Ilett, "Lexus a nexus between cars and phone viruses?", on *CNET*, 26 January 2005, at http://news.com.com/Lexus+a+nexus+between+cars+and+phone+viruses/2100-7349_3-5551367.html

[6] Ross Anderson, "Undermining data privacy in health information", *British Medical Journal* 24 February 2001 (v 322) pp 442-443, at http://bmj.bmjjournals.com/cgi/content/full/322/7284/442

[7] Hal Varian, "New chips can keep a tight rein on consumers, even after they buy a product", *New York Times*, 4 July 2002, at http://www.sims.berkeley.edu/˜hal/people/hal/NYTimes/2002-07-04.html

[8] Ross Anderson, "The Draft IPR Enforcement Directive – A Threat to Competition and to Liberty", Foundation for Information Policy Research, at http://www.fipr.org/copyright/draft-ipr-enforce.html

[9] Ross Anderson, "Cryptography and Competition Policy – Issues with 'Trusted Computing', Workshop on Economics and Information Security 2003, at http://www.cl.cam.ac.uk/ftp/users/rja14/tcpa.pdf

[10] Electronic Security in Cars, http://escar.crypto.rub.de/